

# **FIRE WHITE PAPER, 2009**

## **FIRE: Future Internet Research and Experimentation**

### **Purpose**

The purpose of this White Paper is to define the current framework and status of the FIRE initiative, including ongoing activities and testbeds<sup>1</sup>, also in view of the next call under objective 1.6 of the ICT workprogramme, related to FIRE.

### **Structure**

This Expert Group Report begins (Part A) with a summary of the main aspects of FIRE. i.e:

- The FIRE vision and a description of the (evolving) FIRE initiative;
- a portfolio analysis of the existing FIRE (STREP) projects comparing their coverage also with respect to other ICT research projects and testbeds;
- a positioning of FIRE in the context of ICT Challenge 1;
- relationships with other initiatives
- examples of large-scale experimental research Use Cases

There then follows (Part B) background information, such as more details about the existing projects, potential research activities, open issues, further Use Cases.

The Annex includes the text of the forthcoming call, project descriptions and additional considerations.

---

<sup>1</sup> Whilst the project FEDERICA does not formally belong to the FIRE Unit (F4) of the EC, its goals are aligned closely with those of FIRE, and it is included as a FIRE testbed for the purposes of this document

# Table of Contents

<b>Part A: A Summary of the FIRE Initiative .....</b>	<b>4</b>
<b>1. The FIRE Vision .....</b>	<b>4</b>
1.1 Visionary multidisciplinary Future Internet research defining the Challenges for the FIRE facility .....	4
<i>The concept of experimentally-driven long-term research</i> .....	4
<i>Multidisciplinary approach</i> .....	5
<i>Clean-slate versus incremental approaches</i> .....	5
<i>Research content: open, bottom-up</i> .....	5
1.2 Building a sustainable, dynamic, evolving, large scale FIRE experimentation facility .....	6
<i>The concept of the FIRE facility</i> .....	6
<i>The concept of Federation</i> .....	6
<i>The FIRE facility as a Service Provider</i> .....	7
<b>2. FIRE and its role in Challenge 1 of the ICT Programme .....</b>	<b>8</b>
<b>3. The existing FIRE projects.....</b>	<b>8</b>
3.1 A preliminary portfolio analysis .....	9
3.2 Summary of topics covered by the FIRE, SAC and relevant Future Networks projects .....	11
3.3 FIRE Research projects (STREPs) and their testbed requirements .....	12
3.4 Testbed (STREP and IP) projects with facilities to offer: .....	14
<b>4. Relation of FIRE to Member States and International Initiatives .....</b>	<b>17</b>
<b>5. Examples of Use Cases of large-scale experimental research.....</b>	<b>19</b>
<b>Part B: Background Information .....</b>	<b>22</b>
B.1 FIRE from a Services and Applications Perspective or “The User Perspective” .....	22
B.1.1 User driven innovation .....	23
B.2 The Industry Perspective .....	24
B.2.1 The diversity of objectives in federated testing environments.....	24
B.2.2 Multi-disciplinary aspects .....	24
B.2.3 Increasing the pace of evolution.....	25
B.2.4 SMEs .....	25
B.3 The Research and Academic Perspective.....	25
B.4 FIRE from a Cross Layer / Cross Technology Domain Testbed Perspective .....	26
B.5 FIRE from a socio-economic perspective .....	26
B.6 Relative positioning of the testbed projects .....	27
B.7 FIRE and visionary research on the Future Internet.....	29
B.7.1 Cross-layer, multi-level approach .....	30
B.7.2 Content-centric architectures.....	30
B.7.3 Mobility .....	31
B.7.4 Service Research on FIRE facilities .....	31
B.7.5 Resiliency .....	31
B.7.6 Social networking.....	32
B.7.7 Socio-economic aspects to be strengthened .....	32
B.7.8 Technological issues (non-exhaustive list): .....	32
B.8 The FIRE facility.....	32
B.8.1 Next steps .....	35
B.9 The strong relation between FIRE Research and the FIRE Facility .....	36
B.10 Encourage the re-use of already-developed software/protocol stacks .....	36

B.11 Additional FIRE Use Cases and scenarios .....36

    B.11.1 Cloud Computing (using OneLab2) .....36

    B.11.2 Cloud Computing (federating OneLab2 and Grid 5000) .....37

    B.11.3 Emulating AS Relationships (using FEDERICA).....38

    B.11.4 Public safety .....39

    B.11.5 Mapping Systems for Future Internet Routing (using G-Lab) .....39

    B.11.6 Locality of P2P traffic - University of Würzburg (using G-Lab).....40

    B.11.7 Congestion aware routing and network navigation – A Use Case from the Computer  
 Research Institute of Hungarian Academy of Sciences .....41

    B.11.8 SOA Scenario for Engagement with FIRE Testbeds .....41

    B.11.9 Media content delivery service.....42

    B.11.10 Adaptive monitoring - Fokus (using OneLab2) .....43

    B.11.11 Congestion control - Fokus (using OneLab2) .....43

    B.11.12 Outsourcing of testbed environments - (using PII tools, including Teagle) .....44

    B.11.13 Test of services and applications - (using PII tools) .....44

    B.11.14 Interoperability - (using PII).....44

    B.11.15 Re-use of test suites - (using PII) .....44

    B.11.16 Certification - (using PII) .....45

    B.11.17 Service Oriented Architecture validation - (using SOA4All) .....46

    B.11.18 Congestion control - (using Self-Net) .....47

    B.11.19 Context based services - (using C-CAST) .....47

## **Part A: A Summary of the FIRE Initiative**

### **1. 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.

It covers two strictly interrelated dimensions:

- promoting the investigation of new and possibly unexpected Internet concepts, under an experimentally-driven, long-term yet empirical approach considering the system complexity of the Internet architecture;
- making available to researchers, in industry and academia, an integrated, sustainable, dynamic, large scale experimental facility on future Internet technologies, federating or integrating existing or planned relevant testbeds for emerging or Future Internet technologies.

FIRE is also an excellent environment for doing a rigorous and quantitative research on the broader impact of changes to the Internet, not only in technical but also in socio-economic terms. 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 on social acceptability, user experience, economic viability, etc. and to derive recommendations.

The FIRE activities are strategic for Europe, given their interest both for industry and academia. FIRE is expected to strengthen European industry by providing the innovative concepts and practical 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.

The ultimate benefit expected from FIRE is that EU industry and research are better positioned on Future Internet technologies and services than they were for the first generation. Being successful, FIRE can provide 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. It will also contribute to the development of Academic knowledge and competence in the field. In so doing, the FIRE vision is fully in line with the objectives set in Lisbon.

FIRE is ultimately a mechanism to mobilise European resources around a common theme, creating a critical mass of expertise and encouraging Europeans to work together.

#### **1.1 Visionary multidisciplinary Future Internet research defining the Challenges for the FIRE facility**

##### *The concept of experimentally-driven long-term research*

Large or small research projects in this area should base their research theories and results on the results of experimentations in testbeds. Building associated testbeds and later on providing them within the FIRE facilities can be in the scope of the research projects themselves, or the projects can envisage exploiting existing or planned testbeds, which are already 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 research.

### *Multidisciplinary approach*

Real innovation often comes at the intersection of different disciplines. Moreover, the Internet is a complex system, depending on a multifaceted interrelation between technologies, users, services and applications. Carefully evaluating these interrelations is key to harnessing and exploiting the full potential of the Future Internet for the 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 or even noise in the system to obtain extreme robustness and reliability, for instance in unpredictable environments.
- To explore means by which networks can change, learn and develop spontaneously.
- To apply concepts from the dynamic process of optimisation involved in neural development to the paradigmatic exploration of concepts for the new Internet.
- 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.

### *Clean-slate versus incremental approaches*

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 improvements, the latter implies a radical redesign to offer new abstractions. The latter approach is more likely to be disruptive in terms of technologies, services or business models, and may not preserve “backward compatibility”. Whereas FIRE experimental facilities are obviously meant to be open to both types of approaches, the research carried out under FIRE addresses mainly more visionary though not necessarily disruptive approaches; more risky but at the same time more likely to yield significant improvements.

It must 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.

### *Research content: open, bottom-up*

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 levels of the value chain, from the network connectivity up to service architectures - and, of course, explore different paradigms which might not be based on layered models.

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.

## **1.2 Building a sustainable, dynamic, evolving, large scale FIRE experimentation facility**

### *The concept of the FIRE facility*

The FIRE experimental facility aims to become 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 pre-service trials, which can be used by both industry and academia.

The FIRE experimental facility will be built by gradually federating and integrating existing and new testbeds<sup>2</sup> 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.

### *The concept of Federation*

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. This federation, however, is created with a set of clearly identified objectives in mind, making the federation useful to begin with. Such objectives can be achieving scale, access to heterogeneous resources, open exchange of data, testing of new pre-commercial products, interoperability testing of hardware and software component. It is important to note that not all of these objectives can be realized in a single federation since they can be contradicting (for instance, open exchange of data resulting from pre-commercial testing may be subject to restrictions). Hence, there can be a need for several federation concepts, depending on their set of objectives (and the constraints in achieving them).

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.

While a completely integrated facility may be desirable, this will be difficult to achieve both in terms of cost and of time horizon. Moreover, a single form of federation can be impossible, due to conflicting test requirements or commercial restrictions. Hence, the establishment of different federations must be considered, e.g. one that targets early stage (open) research exchange to stimulate

---

<sup>2</sup> A testbed here is to be understood as an environment allowing rigorous, transparent and replicable testing and experimentation in research and development projects

new and innovative ideas, another one targeting pre-commercial exploitation (including the generation of patent and the alignment with upcoming standards).

### *Key challenges for the FIRE facility*

Some key challenges for the FIRE facility are:

- Federation
- Governance of the federation
- Incentivising the federation
- Sustaining the federation
- Sustaining the measurement data
- Virtualization
- Monitoring
- Reproducibility
- Benchmarking
- Security
- Economics of the facility
- Legal responsibilities and liabilities, IPR, etc.
- Risks
- Aligning offer to demand
- Overcome problems of timing (facility projects may end before they can be used)
- Customer-friendliness

### *The FIRE facility as a Service Provider*

In order for the testbeds to be further developed and evolve over time, it will be of major importance that they are used, especially by the FIRE research projects and other research projects under Challenge 1. Only then will it be possible for the testbed providers to understand the requirements of testbed users of different types. Therefore, for the providers, encouragement or some kind of enforcement of the testbed usage would make sense.

On the other hand, it is well understood that using the testbeds might not be easy for all users because of a lack of knowledge of the capabilities and how to access them. The testbed providers will therefore produce information packs for potential users, and users will write Use Cases to define the experiments they want to perform on the testbeds.

In order to progress the work on how such "service provision" might be done, a small working group has been established to draw a map of existing facility prototypes and their relations and to derive from there an outline of the architectural principles for a high-level FIRE federation architecture for federating/ integrating the evolving FIRE facility prototypes and related projects and infrastructures, addressing not only technical but also operational and legal aspects as well as issues related to different business models. Through this, the group shall set the framework for the federation/ integration of future facility projects submitted under the next FIRE Calls. The basic idea is to get a modular framework of largely independent components (projects) with defined and agreed common services and interfaces/gateways. This working group is considered as the beginning of an integration process which obviously will require refinement and implementation by the running and coming FIRE Facility projects.

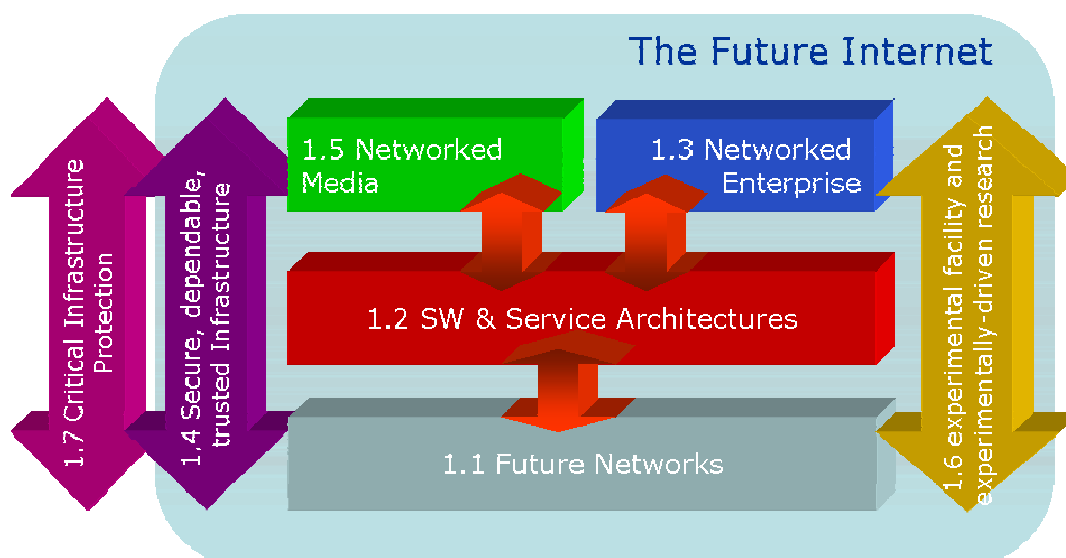
## 2. FIRE and its role in Challenge 1 of the ICT Programme

FIRE is open to include testbeds from other projects of Challenge 1 and beyond. Additionally, projects (including broad system-level type projects) from other parts of the Framework Programme should 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 making socio-economic impact assessments of future changes to the Internet. However, the FIRE research projects themselves should aim at becoming the most demanding users of the FIRE facility.

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

- The integration of projects in a sustainable environment for investigation and experimentation of new (evolutionary and revolutionary) paradigms.
- The empirical, experimentally-driven character of the research (not just paperwork). The research must be experimentally-driven including large scale testing.
- The bottom-up approach, multidisciplinary, considering the internet as a complex system. The research must cut across objectives of Challenge 1 and must take a broad system perspective.
- The strong attention to the social, economic and environmental impacts of the proposed research.
- The research activities must exploit the FIRE experimental facility as much as possible - missing testbed components may be provided for federation within FIRE. These projects are the "hard" test and Use Cases for the facility.
- The experimental facility projects must have the potential to provide a service to others in the future (as a testbed or part of one).

The positioning of the FIRE activities in Challenge 1 (from the perspective of the FP7 organisation):



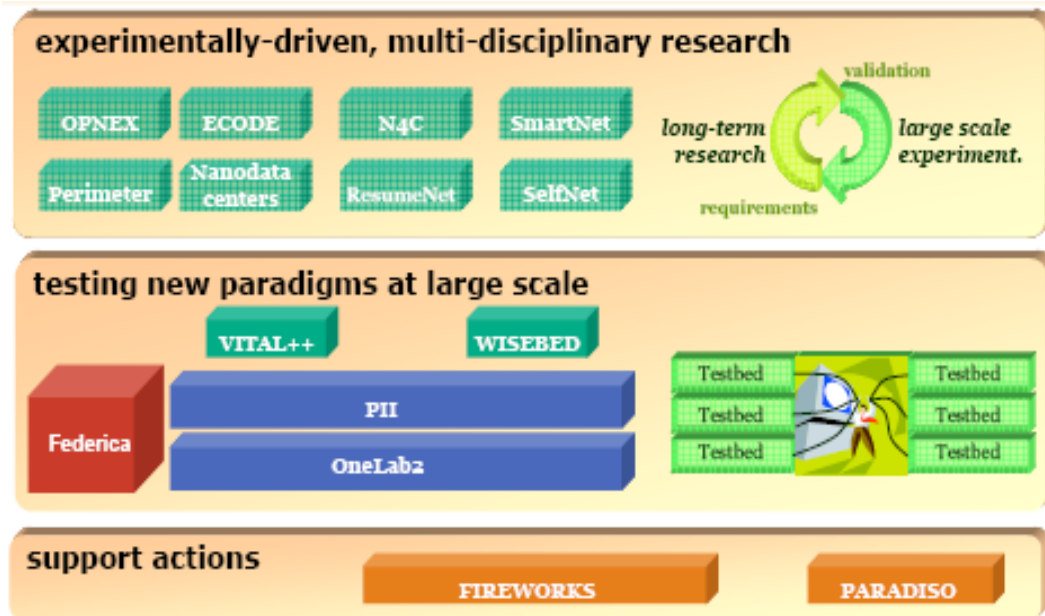
## 3. The existing FIRE projects

In the FP7 ICT Call 2, for Objective 1.6 “Future Internet experimental facility and experimentally-driven research” 10 STREPS (8 for research and 2 for testbeds), 2 IPs (testbeds) and 2 accompanying measures were selected for funding.

- STREP projects: researching into new protocols and paradigms;

- STREP and IP projects: building testbeds and the conditions for their use.
- The two coordination and support actions look at the field from a strategy and trend-analysing viewpoint, enabling multi-disciplinary discussion on prospects of FIRE, also beyond Europe.

FIRE projects can be clustered in the following way:



### 3.1 A preliminary portfolio analysis

ECODE, NANODATACENTERS, N4C, OPNEX and Vital++ deliver novel architectures and techniques for **routing and transferring data**. ECODE and OPNEX work on routing, whereas NANODATACENTERS, N4C and Vital++ have opportunistic networking in common.

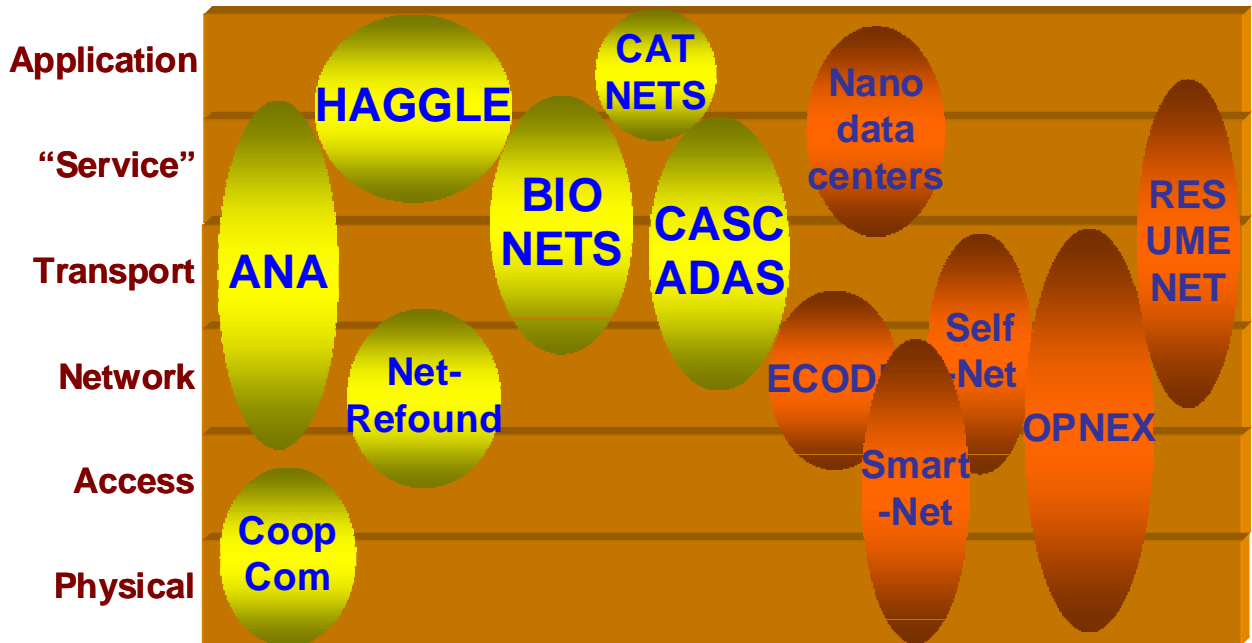
PERIMETER, ResumeNet, Self-NET and SMART-Net are studying networks or their parts at a conceptual level and are proposing new **management methods** for them. PERIMETER and SMART-Net are concentrating on the access and connectivity issues, ResumeNet induces resilience through the layers and Self-NET delivers a framework for self-managed network elements.

The **testbed** facility projects target the establishment of federations, each under a different set of objectives, requiring complementary approaches in achieving the federations. These different objectives are, for instance, reflected in the desired scale, the openness of the facilities, the extent of data sharing, the integration of local policies and/or the user bases and therefore user profiles for each federation. More concretely:

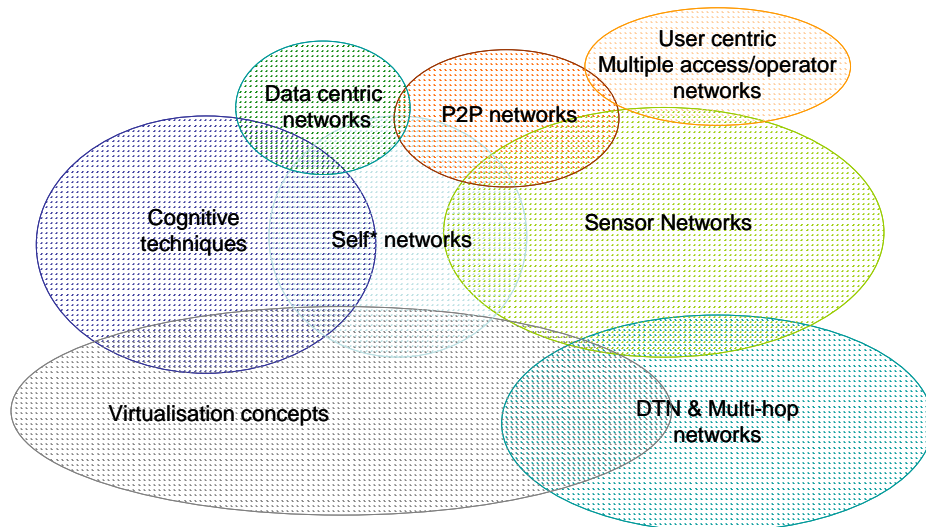
OneLab2 builds on an existing open global federation of networks and devices for networking experimentation that addresses the scientific research community, whereas PII builds up pilot federations of technologies and systems, which can be used for testing (for example) interoperability issues and issues closer to the market. In parallel they do research on a universal brokering tool for testbed services. WISEBED builds wireless sensor testbed federation and Vital++ interconnects IMS testbeds, experimenting with peer-to-peer technology. N4C experiments with DTN methodologies based on a concept developed in the Swedish Lapland. PERIMETER has two interconnected large-

scale testbeds with user groups validating the paradigm for seamless connectivity that the project develops. NANODATACENTERS contributes to OneLab2 by bringing in its validation environment for the distributed data storage and transfer architecture as a part of the OneLab2 federation. The rest of the STREP projects will be customers of the FIRE testing facilities.

Self-Management is addressed in several FIRE projects. The following picture shows how different FIRE and SAC projects address self-management aspects at different OSI layers.



The following picture shows instead the area coverage of the current FIRE research testbeds



### 3.2 Summary of topics covered by the FIRE, SAC and relevant Future Networks projects

	ECODE	N4C	NADA	OPNEX	PERIMETER	ResumeNet	Self-NET	SMART-Net	OneLab2	PII	WISEBED	Vital++	Federica	BIONETS	ANA	HAGGLE	CASCADAS	4WARD	TRILOGY	PSIRP	CHIANTI	AUTOI	MOMENT	N-CRAVE	SmoothIT	EFIPSANS	Euro-NF
Autonomic Configuration	X			X	X	X	X							X	X	X	X	X				X				X	
Best Connectivity and Quality					X								X						X		X				X		X
Control Plane Functionality						X	X	X		X		X							X	X					X		X
DT, Opportunistic Networking		X							X					X		X		X			X						
Data Centric Networking			X						X									X		X							
Passive and Active Measurement Technologies	X						X		X				X										X				
Incentive mechanisms for sharing resources									X	X	X									X					X		
Multi-hop ad hoc wireless sensor networks and mesh networks optimisation				X			X	X	X		X																X
Cognitive techniques and components	X					X	X							X		X	X										
P2P architectures			X									X		X		X				X					X		
Virtualisation of resources			X						X	X			X					X									X
Optimisation frameworks				X																							
Novel Architectural principles	X		X	X		X	X							X	X	X	X	X		X							X
Resiliency and Fault Tolerance						X								X			X										
Cross-layer approaches				X			X	X			X			X		X											
User-centric design					X											X											
Monitoring and Self Governance														X	X	X	X	X				X				X	
Network Coding														X										X			

### 3.3 FIRE Research projects (STREPs) and their testbed requirements

Project	ECODE	N4C	NANODATACENTERS	OPNEX
<b>Main topics</b>	<p><b>Cognitive routing system using machine learning techniques:</b></p> <ul style="list-style-type: none"> <li>- Combining passive and active measurement for distributed detection of traffic anomalies</li> <li>- Efficient path ranking based on QoS</li> <li>- Inter-domain routing system dynamics</li> <li>- Autonomic configuration</li> </ul>	<p><b>Opportunistic networking architecture</b> for `communication challenged` communities <b>using DTN</b></p>	<p><b>Nano datacenters using peer-to-peer communication infrastructure:</b></p> <ul style="list-style-type: none"> <li>- Utilising the capacity at the edges of the network instead of growing centralised data centres</li> <li>- Using virtualisation for partitioning and management of resources</li> </ul>	<p><b>System optimisation for multi-hop network design</b> taking advantage of:</p> <ul style="list-style-type: none"> <li>- Backpressure principle</li> <li>- Max-weight scheduling</li> <li>- Utility optimisation congestion control</li> <li>- Primal-dual method</li> <li>- Autonomic configuration</li> </ul>
<b>Outcome</b>	<p>Additional cognitive component for routing design adopted from machine learning</p> <p>Revisited loosely coupled control infrastructure for the Internet</p> <p>New <i>online</i> and distributed learning algorithms</p>	<p>Development of the DTN bundle protocol and appropriate routing technology to transfer messages through opportunistic encounters between nodes</p> <p>Adaptation of applications to suit the DTN environment</p> <p>Sustainable testbed for ONA</p>	<p>Disruptive service and content delivery architecture using edge computing, with:</p> <ul style="list-style-type: none"> <li>- security and incentive mechanisms</li> <li>- cost effectiveness</li> <li>- content near to the user (QoS, cost, etc.)</li> </ul>	<p>New protocol stack for multi-hop wireless network</p> <p>To extract full transport capacity by converting algorithms to protocols and architectures</p> <p>Decentralised low complexity, low-overhead implementation</p>
<b>Facility needs / outputs</b>	<p>Need:</p> <ul style="list-style-type: none"> <li>- Virtual and physical: e.g, OneLab and iLab.t</li> </ul>	<p>Output:</p> <ul style="list-style-type: none"> <li>- Extension of the SNC architecture for wider usage together with applications adapted for use in the DTN environment.</li> </ul> <p>Need:</p> <ul style="list-style-type: none"> <li>- Seamless mobility between DTN supported regions and the conventional Internet</li> </ul>	<p>Output:</p> <ul style="list-style-type: none"> <li>- IPTV/VoD server connected to a set of P2P gateways in OneLab2</li> </ul>	<p>Need:</p> <ul style="list-style-type: none"> <li>- Wireless testbed (802.11) and one with customised sensor nodes</li> </ul>

Project	PERIMETER	ResumeNet	Self-NET	SMART-Net
<b>Main topics</b>	<p><b>Paradigm for user-centric, seamless mobile communications</b></p> <ul style="list-style-type: none"> <li>- QoE concept with A3M and AAA</li> <li>- ABC (Always Best Connected) in multiple access multiple operators network</li> </ul>	<p><b>Architectural approach for Internet resilience</b> that is multilevel, systemic and systematic including:</p> <ul style="list-style-type: none"> <li>- Protocol layer</li> <li>- Data, control and management planes</li> <li>- Fault-tolerant network components from survivable network topologies (Internet including end-systems)</li> </ul>	<p><b>Cognitive self-managed elements for network management</b> based on:</p> <ul style="list-style-type: none"> <li>- Distributed Cognitive Cycle for System Management</li> <li>- Cognitive behaviour</li> <li>- High degree of autonomy</li> <li>- Localised management decisions</li> <li>- Peer-to-peer style distribution of responsibilities</li> </ul>	<p><b>Architecture for next generation wireless mesh networks by deploying smart antennas</b></p> <ul style="list-style-type: none"> <li>- Access infrastructure distributed across the network</li> <li>- Smart antennas for the use with each subscriber unit</li> <li>- Multi-RAT operation</li> <li>- Autonomic configuration</li> </ul>
<b>Outcome</b>	<p>Middleware architecture to support:</p> <ul style="list-style-type: none"> <li>- Generic QoE models</li> <li>- Signalling-content adaptation</li> <li>- Prototype applications and services</li> </ul>	<p>Resilient framework for networking by defining the relationship between:</p> <ul style="list-style-type: none"> <li>- Self-protection</li> <li>- Redundancy</li> <li>- Diversity</li> <li>- Resource tradeoffs</li> </ul> <p>High-level model applied to e.g:</p> <ul style="list-style-type: none"> <li>- Routing</li> <li>- E2e protocols</li> </ul>	<p>A framework for management actions and protocol stack design:</p> <ul style="list-style-type: none"> <li>- flexible and dynamic schemes</li> <li>- autonomic models</li> <li>- localised management decisions</li> </ul>	<p>Advanced protocols and procedures for SMART-Net architecture</p> <ul style="list-style-type: none"> <li>- Scalability</li> <li>- Reliability</li> <li>- Security</li> </ul> <p>Comprehensive and configurable experimental facilities for wireless mesh networks</p>
<b>Facility need/output</b>	<p>Outcome (pilot):</p> <ul style="list-style-type: none"> <li>- Interconnected large-scale testbeds of real users (mobile networks, user communities)</li> </ul>	<p>Need:</p> <ul style="list-style-type: none"> <li>- Tbd (on top of existing to-be-enhanced testbeds)</li> </ul>	<p>Outcome:</p> <ul style="list-style-type: none"> <li>- A prototyping environment for Self-NET solutions</li> </ul>	<p>Outcome/Need:</p> <ul style="list-style-type: none"> <li>- Large-scale experimental environment (SMART-Net) interconnecting:                             <ul style="list-style-type: none"> <li>- Virtual testbeds</li> <li>- Real testbeds</li> </ul> </li> </ul>

**3.4 Testbed (STREP and IP) projects with facilities to offer:**

	<b>OneLab2</b>	<b>PII</b>	<b>Vital++</b>	<b>WISEBED</b>	<b>FEDERICA</b>
<b>Application domains</b>	Typically scientific / research applications (e.g. IP networking, distributed systems) in a real world environment	Typically pre-service trials of commercial applications. Industry focus	Peer-to-peer content distribution, combined with IMS-like Control Plane. Typically new services and content delivery mechanisms for telecom / Internet service and network environments	Typically research applications in distributed wireless sensor networking environments (e.g. status of machines, environment, food-safety, medical status, energy management, building automation)	Any application not requiring specialized HW. The focus is on new networking services in Layer 2 and 3 (eg. routing, protocol interoperability, virtualisation). Access to the control of Layer 1 services in a later phase
<b>Application installation</b>	Users install their applications on their own site	Users may install their applications on their own site, or bring to one of the PII testbeds		Users install their applications at any of the WISEBED test laboratories across Europe	Users may install their applications on their own site, or at FEDERICA nodes and connect using Internet or a dedicated circuit (to be agreed)
<b>Service Delivery Platforms / Service Overlays</b>	Users control their own experiments	Converging network, service platform and application infrastructures (SOA). NGN, TISPAN, IMS. PII has complete control over the resources	IMS		Control and configuration of slices is via a Web-based GUI. The user can install its own management system
<b>Networks</b>	Both public (the general Internet) and private	Fixed (xDSL, CATV) and Wireless (3G, LTE)	Interconnected IMS testbeds across Europe	Wireless sensor network test laboratories across Europe, interconnected over IP	Up to 1 Gps set of Ethernet circuits created using GFP on GEANT. Lower speeds are available through virtualisation (MPLS or VLAN)
<b>Access Connectivity</b>	If using the general	Access is via controlled	3G		Traffic can be injected

	Internet (shared resources), experiments will not yield exact reproducible results	telco networks. Experiments will be generally reproducible (bearing in mind the lack of strict guarantees on any Internet segments)			from anywhere with access to the Internet (but resources and QoS are not guaranteed outside the FEDERICA network).
<b>End Systems / Devices</b>	Hundreds of (Linux) PCs	Enhanced versions of existing terminal types			Workstations with virtualization software
<b>User control and management</b>		Not applicable			The user is given the capability to self-manage the resources within his/her slice, down to the level of transport and routing protocols.
<b>Security</b>					The virtual slices are completely independent from each other. Security (encryption) of the data within a slice is a matter for the user
<b>Ad hoc communications</b>		Not supported			No direct support, but open for the attachment of ad hoc systems and devices
<b>Monitoring</b>	Embedded passive and active measurement technologies	Y		Y (IPv6 is used for the monitoring traffic)	GÉANT tools available for monitoring the status of interfaces and links, and for localising faults. Flow monitoring on selected circuits only.
<b>Hosting users equipment</b>					Possible with some space a constraints (about 1/3 <sup>rd</sup> of a rack in each PoP)
<b>Other</b>	Incentive mechanisms for sharing resources	Incentive mechanisms for sharing resources		Incentive mechanisms for sharing resources	Standardisation of virtualisation (new IETF

				Standardisation of wireless sensor network testbed interconnection	WG)
<b>General testbed issues:</b>					
<b>Conditions for using the testbed (open access)</b>	Free, but users must commit to provide 2 PCs for use by the OneLab2 community	Telco-quality networks are used. A fee will therefore be charged according to the resources used and duration of the experiment		Free, until the end of the project. Afterwards, users will have to pay according to the resources used	Free, but must not be used for carrying commercial traffic. The experimentation period may be limited to 3 months (depends on demand). Physical access must be organised (and paid for) by the user. Results should be shared (if possible) with FEDERICA
<b>Is the testbed scope extensible?</b>	Yes. OneLab2 is the European part of PlanetLab. The scope is limited only by the capability of the network; i.e. the characteristics of the Internet, available bandwidth and computing capabilities of the end devices). wireless networks and DTNs extensions planned FEDERICA federation is being investigated	Yes (open for federation, also with non-IMS platforms)		Yes (extending in scale and diversity in time. Developing more advanced hardware, software and data handling applicable for wireless sensor networks)	The geographical scope is being extended, but will generally follow the GÉANT footprint. The network is open for federation with other testbeds in order to extend the geographical and technological scope. Federation with OneLab2 is being investigated
<b>Are the testbed components and toolkits open source?</b>	Yes	No			Yes (based on VMware and open source developments)

## 4. Relation of FIRE to Member States and International Initiatives

FIRE is among the largest initiatives in Europe and worldwide for building an experimental facility for Future Internet research. After its preparatory phase and the launch of the first wave of FIRE projects nearly a year ago, FIRE has established itself on eye-level with other initiatives, whose prototypes of experimental facilities are at a similar stage and maturity.

FIRE's unique features as compared to other initiatives and projects around the world are:

- **The close link between the FIRE Facility and the FIRE Research**, which pushes the limits and defines the challenges for the FIRE Facility through visionary, multidisciplinary research.
- **The positioning of the FIRE Facility as an infrastructure for Future Internet research in the broadest sense** considering the Internet as a complex system and therefore addressing all the associated aspects of the Internet in a holistic vision, and at all relevant levels and layers including network architectures and technologies, service architectures and platforms, networked media and trustworthy infrastructures.
- **The broad scope of FIRE** ranging from proof-of-concept testbeds primarily addressing the longer term research community to pre-commercial testbeds addressing experimentation closer to the market and being of particular interest for industrial players.

Within FIRE a discussion has started on how to collaborate and whether to eventually federate the different facility prototypes in order to broaden the scope, scale and diversity of the resources offered to the industrial and academic research community. This effort of collaboration and federation cannot stop at the boundaries of FIRE. Instead, mechanisms need to be found to include national resources of the EU Member States and Associated Countries. Discussions are on-going with the most-advanced and relevant national initiatives in order to better explore synergies and potential for cooperation at different levels. It must be recognized though, that the objectives of a federation of national resources are unlikely to be realized in a single form of federation since they can or are contradicting and are therefore hard to align (common objectives are a cornerstone of ANY federation). Hence, at least as a starting point, the establishment of different federations must be considered between facilities with at least some shared common objectives. As the Future Internet is a global issue, collaboration on an international scale is a must. First collaboration and federation between initiatives and projects on the EU side and in the US, Japan, and Korea has started.

In an Annex, a list of EU national and international activities related to FIRE is provided.

In Europe, during the Future Internet Assembly, the FIRE launch, and the FIRE Expert Group Meetings, a dialogue with many Member States Future Internet initiatives and projects has started. The discussions with the German and the French initiatives are most advanced:

- German initiative G-Lab consists of a research and experimental facility used to investigate the interplay between new technologies and the requirements of emerging applications for the Future Internet. G-Lab partners develop a secure and reliable platform for applications and services with special emphasis on new mechanisms for routing, addressing, etc., and their investigation with regard to practical feasibility, scalability, and performance in realistic environments. In their second phase G-Lab plans to intensify its collaboration with industry at the relevant levels of the value chain. G-Lab and Onelab2 are discussing a slice-based federation federation for testing on larger scale, testing algorithms in a heterogeneous environment, testing beyond one Tier, testing under more realistic conditions, and for extending the geographic distance between the sites (for example to get longer delays).

- French project ALADDIN-Grid'5000 is an initiative aiming at developing a highly reconfigurable, controllable and monitorable experimental large-scale distributed platform. Although, its design targets experiments in the area of Grids, its high degree of reconfigurability makes it suitable to carry out experiments in the area of service infrastructures, cloud computing and the Internet of Servers. Due to their complementarities, ALADDIN-Grid'5000 and Onelab2 are considering federation of both facilities driven by a common use case (Section n.m). Such a vertical federation would allow experimentation across many layers of the internet in an integrated fashion, thereby bringing together the networking and the services research communities.

Future Internet initiatives in several other Member States and Associated Countries have a strong focus on experimentation. Here only an incomplete list of some highlights can be included:

- The Belgian project ilab.t of IBBT supports experimental multidisciplinary research for FI technologies and innovative ICT services and applications by industry and academia. Among others they focus on wireless technologies, and a "virtual wall". Ilab.t is part of ECODE and is collaborating with their IBBT Living Lab towards bringing end-users into the experiment.
- In Italy, the National Research Council (CNR) features several large-scale projects on Future Internet topics ranging from integration of enabling technologies to management of multimedia content, from global security and public safety to grid services and cloud computing, personalised smart data mining and knowledge extraction techniques. Discussions are ongoing about how to exploit synergies among these initiatives and the EU Future Internet activities. The National Research and Education Network (GARR) coordinated the FEDERICA project and may act as a facilitator.
  - RedIRIS, the Spanish academic and research network, actively participates in the FEDERICA project, on aspects like user groups, network infrastructure design, authentication and authorisation, and new resource management models. Spanish funded project Project PASITO (Telecommunications Service Analysis Platform) aims at the deployment of an experimentation platform for new Internet services and protocols as well as the provisioning of an Open platform in which other academic and research centre research groups and, occasionally, operators and manufacturers, collaborate.
- The Finnish ICT SHOK initiative has a strong focus on wireless technologies and on user-driven open innovation bringing the user into the loop through experimentation at all stages of the technology development and adoption cycle. Finnish partners are bringing this experience to a European scale through their strong involvement in FIRE project PII.
- In Sweden several national projects with a strong experimentation dimension are supported, most of them having a strong application dimension like IPTV, interactive TV, user behavior services, just to name some. Projects often involve industrial and societal players at all levels of the value chain, some of them bringing users into the experiment at a broad scale.
- In the Netherlands, testbeds focus on wide-area distributed systems, lambda networking as well as photonics-based Grids for very high quality digital media.
- In Hungary, experimentation is driven by the national research network, which provides experimentation resources. Driven by realistic use-cases, Hungarian organizations are investigating and plan to pilot a potential federation between projects Federica and Onelab2.

Outside Europe, the US and Japan are most active in Future Internet research and in the provisioning of experimental facilities, here just some highlights:

- In its first spiral, the US-GENI initiative supports five to some extent complementary clusters of projects each of them being built on a different control plane, targeting different research user communities, and starting from existing infrastructures:

- DETER (targeted at security research community);
- Planetlab (targeted at distributed systems and networking community),
- ProtoGENI/Emulab (Internet2 backbone, programmable edge clusters); ORBIT/Winlab (targeted at the wireless research communities), and
- ORCA (including compute resource, service-oriented, no baseline project).
- ORBIT Management Framework (OMF).

Led by NSF and the EC, FIRE and GENI regularly organise joint workshops (Dec 2008 in Madrid, July 2009 in Seattle) for bringing key researchers together and exploiting synergies. Due to its common origins in Planetlab, FIRE project Onelab2 is closely collaborating with GENI-Planetlab. Through managing Planetlab Europe, Onelab2 is part of a worldwide federation with Planetlab US (PLC) and Planetlab Japan. Onelab2 is bringing new innovative dimensions into Planetlab, e.g. related to wireless and to autonomic communication. FIRE project PII has recently started to collaborate with GENI-ORCA. Both projects develop a universal brokering tool for testbed services based on similar principles. The projects are complementary with ORCA being strong on addressing the higher level services dimension and with PII addressing industrial aspects and pre-commercial testbeds. Belgian project ilab.t is collaborating with GENI-EMULAB. Federica is working closely with the Internet2 consortium.

- The Japanese Future Internet research landscape is based on two concepts: Next Generation Network (NXGN) and New Generation Network (NWGN), the latter using new paradigms and architectures, with a view to replacing the current IP. For the moment, the AKARI Architecture project works on the design of the future network, supported by the experimental network project JGN2 (Japan Gigabit Network 2) and its successor JGN2plus. Besides at more general EU-Japanese for a on ICT research, the Japanese and the European Future Internet research communities meet regularly (June 2008 in Brussels, July 2009 in Tokyo) to inform each other, to exchange views, and to start or deepen collaboration. On both sides there are similar long term goals in redesigning the Internet. Related to experimentation, EU-Japan collaboration has led to the “decentralisation of PlanetLab” towards a federation of Planetlab US (PLC), Planetlab Europe (operated by Onelab2) and Planetlab Japan.
- During the OECD Conference and the Korean-EU bilateral meeting in Seoul in June 2008 and a bilateral event between Korea and the EU in November 2008 in Brussels, representatives of the FIRE community met with the Korean Future Internet research and experimentation community.

In addition there are close contacts between FIRE partners and organisations in Canada (related to their 0-carbon initiative) and Australia (Onelab2 has an Australian partner who is also part in US-GENI Orbit). First more informal discussions have started as well with Brazil and China.

## 5. Examples of Use Cases of large-scale experimental research

The presented examples showcase in a concrete way some possible uses of the FIRE facility projects and the way experimentally driven research can be performed. These examples are not thoroughly developed as they just provide ideas. and it is hoped that they provide inspiration for similar experimentations. The main reason for using FIRE facility projects in those use cases is to test at greater scale and at a more diverse environment than the one that the individual research projects can afford on their own and get more close to real network conditions. In some other use cases it is important to have a large scale but at the same time controllable environment where parameters can be tuned and credible measurements can be taken. More details can be found in part B of this document.

<b>Use case</b>	<b>What?</b>	<b>How?</b>	<b>Why (FIRE)?</b>	<b>Challenges</b>
Cloud Computing	Design new scalable, self-managed, energy efficient cloud architectures while respecting SLAs.	Combine O2 and Grid'5000	Research in large scale and in close to real conditions offered by the combination	Combine 2 different test beds (O2/ Grid'5000) with their own administrative/security policies and network configurations
Emulating Autonomous Systems (AS) Relationships	Build a virtual topology over the existing Internet, to test the inter-AS relations	Set AS relationships, while conducting experiments of various service and/or networking technology. Then visualize the resulting topology and give a basic economic understanding of the new relationships on FEDERICA	Provide close to real conditions (changing AS relations) to test new inter-domain publish-subscribe solutions, or content delivery, federated sensor solutions, cloud Computing scenarios etc.	Emulate virtual future Internet topologies over the existing Internet
Mapping Systems for Future Internet Routing –	Test a hierarchical routing: current Internet routing does not scale since the number of entries in BGP routing tables	Apply map and encapsulation techniques and use G-Lab to test and improve the mapping system	G-Lab offers controlled environment to test and monitor the mapping system	
Locality of P2P traffic	Test new traffic management mechanisms based on cost and QoS.	Test in large-scale overlay networks including various ISPs and various underlying physical substrates Federate SmoothIT trial with the G-Lab	Validate the results for realistic scenarios and generate realistic traffic pattern within the overlay.	
Congestion aware routing and network navigation	Test various routing strategies about new routing principles reacting to congestion on links, where end users can select specific paths for the delivery of their data packets.	Test new routing and select specific paths reacting to congestion by using federation between FEDERICA and OneLab2	This use case needs large scale testing and is important for industry e.g. IPTV providers, P2P video service providers, fibre to the home network providers, intelligent home applications and large scale Internet of Things	Current test beds prohibit multimedia traffic flooding the core network. The physical access to FEDERICA and pricing of high (10-100 Gbit/s) traffic from 50-100 different

Use case	What?	How?	Why (FIRE)?	Challenges
			(RFID) players.	network locations should be solved.
Service oriented architecture validation	Investigate the interoperation of existing heavyweight and evolving lightweight SOA platforms and their underlying service and networking infrastructures.	Implement concrete use scenarios with real users on each of the platforms in an interoperable way and compare behaviour, ease of us, QoE, ... Assess impact of changes of components in the underlying network and service infrastructure.	Bring different communities (SOA- networking researchers) to work together on a system-level approach for the Future Internet.	Federation of SOA testbeds with Network infrastructure testbeds which could be provided by PII or Onelab2. Include in experimentation also proprietary platforms.
Distributed High-performance computing	Test the scaling of web services in terms of algorithms for distributing the services to the clients under the problem of congestion at limited network speeds.	Federica can guarantee bandwidth between distributed components of a service offering. Protocols that offer increased control over backbone connections can be tested with Onelab2	The combination of those test beds will provide a controlled environment for testing	
Experimental verification of different Internetworking protocols	Experimental verification of Internetworking protocols	Different peering models are designed for a set of Internetworking protocols. Evaluation of the performance and scalability of different Internetworking architectures.	The experimentation will allow estimating the impact of different economic conditions on the viability of various proposed Internetworking solutions.	
Adaptive monitoring	Monitoring infrastructure for self-managing autonomic networks.	Adaptive and programmable monitoring with measurement functions placed and configured dynamically in the network.	The usage of OneLab2 testbed allows for validation and practical experimentation in real-world like environment..	A testbed with nodes based on the ANAcCore is planned for PlanetLab Europe
Congestion control	Congestion Management of WiMAX broadband access provision to consumers.	Consumers are using normal In case of congestion on the WiMAX wireless link the bandwidth for each flow is decreased in order to maintain the current number of voice calls with a voice codec featuring higher compression..	The experimental approach is necessary to perform reconfiguration resulting in an overall decrease of traffic over the previously congested wireless links while preserving Quality of Service	

## Part B: Background Information

This part of the Expert Group Report gives additional background information, such as potential research activities, the FIRE constituency, strategic issues, further Use Cases and relationships to national and international initiatives, and more details about the existing projects.

### B.1 FIRE from a Services and Applications Perspective or “The User Perspective”

In general, large-scale experimental testbeds are needed to incubate emerging architectural and operational paradigms as a precursor to any commercial deployment. For that purpose, testbeds need to enable the investigation of properties at all levels ranging from real-world applications to service architectures down to the physical level of the underlying infrastructure. The testing of applications is therefore an important component of the design of such networks – on the other hand, end user applications should be innovative and designed to thoroughly explore the testbed capabilities.

In general, potential users find it difficult to access the current facilities (either in terms of physical connectivity or knowing what hosting/experimentation/control/monitoring and measurement facilities are available) and both parties need to invest time and effort in order to bridge this gap.

A more non-technological view on experimental research raises the question as to what role end users will play in Future Internet research. Although not every experiment must involve end users, (e.g. a component-level stress testing of some new forwarding node does not require any exposure to end users), user-centric experimentation will become increasingly commonplace. While concepts like LivingLabs aim at providing such user-centric experimentation facility, it is the proper link to state-of-the-art component and systems research that needs to increase. In other words, there needs to be a stronger emphasis in Future Internet research to seek the involvement of - and test with - end users, either on the co-development or usage level. This becomes particularly important when targeting economic or societal questions with respect to certain new fields of technologies.

Such growing involvement of end users also raises the question as to how we can integrate new forms of innovation in research with innovation in products and services within the deployed ‘current Internet’.

Users should be involved in the development of information systems for several reasons. In the following a selection of these reasons and benefits is presented and clustered into three end-motives: ethical (democracy), curiosity (theoretical), and economic (pragmatic) (Bergvall-Kåreborn and Ståhlbröst 2008):

- *Ethical* (democracy) perspective: According to the ethical motive for user participation, people have a moral right to influence their own destiny, and users have a right to influence technological decisions affecting their private and professional life. To do this they need to participate in the design process and be given influence and mandate in the decision-making process. The guiding concepts for this perspective are democracy, power to the people, and improved quality of life. Hence, the focus is on strengthening the position of the users of technology (Bergvall-Kåreborn and Ståhlbröst 2008).
- *Curiosity* (theoretical) perspective: The curiosity motive for involving users is to learn more about the nature of participation, and operational findings from this type of research can be used both for

ethical and pragmatic motives. This involves exploring issues such as the location of knowledge and knowledge sharing, different degrees and types of participation, and the ways in which corporate and national culture affect participation. The guiding concepts for this perspective are cooperation, communication, and mutual learning between participants and contributors. Hence, the focus is on creating gains for all key stakeholders (Bergvall-Kåreborn and Ståhlbröst 2008).

- *Economic* (pragmatic) perspective: The main motive for user participation from an economic perspective is to get the job done better. This usually entails improving innovation quality (effectiveness and efficiency) and the acceptance of the innovation. When it comes to the quality of the innovation, the emphasis is on improving functional requirements by increasing designers understanding of the actual use context and gain access to users knowledge, skills, and expertise. When it comes to system acceptance, commitment and buy-in among the users are crucial. The key concept for this perspective is getting the job done better and more cost-effectively.

### B.1.1 User driven innovation

End users may push innovations into devices by demanding application and features based on their real demand, wishes or experiences. One new opportunity in relation to efficient development of mobile and autonomous real time services is the possibility for people in general, both professional and private, to develop their own services where and when they experience a need of IT-support.

There are two kinds of users:

- users of the testbeds, and
- end users i.e. the users of the technology, products, services that the users are testing on the testbed

Focusing on the end users (e.g. existing Living Labs), some open issues can be identified:

- Is the PII Teagle tool a suitable interface for connecting the FIRE facility to the end user? i.e. will it have simple "click & tick" boxes to discover and add other end users ("co-creators"), and configure the testbed(s) according to what the (e.g. LivingLab) end users want for their tests?
- Can the PII Teagle tool also be used to est social and economic dimensions, e.g. how end users can do business between themselves and with the FIRE testbeds; whether federated, or individual?
- End users can be considered as an asset to the FIRE eco system, and it should be made easy for them to be reached and easy and attractive for them to use the FIRE facility. This leads to the question, how users could and should be rewarded for their participation.
- Experiment repeatability is important. However, whilst changes should not cause unexpected technical problems, it is also important to know if the last change/adjustment was positive or negative for the end user. This, however, requires a degree of reproducibility that often constrains the spectrum of tests that can be done.
- Can end-users do what they want to do with the product, service?
- Do they want the product, service?
- Can they afford the product, service?
- How can users be motivated to participate in tests and how can they be located and encouraged to participate?

These issues will be among those discussed at the "FIREweek & LivingLabs" event in Luleå (1<sup>st</sup> - 3<sup>rd</sup> July, 2009).

## B.2 The Industry Perspective

This sub-section highlights aspects that are seen to be crucial from an industrial aspect. The text does not claim to be exhaustive in its sets of arguments nor representative with regards to the breadth of consulted industrial representatives. But it does claim to strike a chord that is relevant for the FIRE discussions.

### B.2.1 The diversity of objectives in federated testing environments

FIRE is driven by the concept of *federating* various testbeds and encouraging experimentally-driven research to make proper use of these growing offers of federations in their quest to drive forward the development of the Future Internet.

It is well understood that the very notion of a *federation* requires a common set of objectives to be present when constructing the federation itself. Industry recognizes however that the objectives of federations can be contradictory although being perfectly suitable and even desirable on their own. For instance, the objective of a federation to improve scientific outcome through openness and community creation can be at odds with the insertion and usage of proprietary (corporate) information within a federation.

Another important aspect of objectives and their potential contradiction comes from the proximity to 'real systems'. While component-level testing is often striving for reproducibility as a crucial objective, such reproducibility becomes more relaxed as a requirement when considering a system-level set of experiments, recognizing the difficulty to reset all crucial parameters to the same initial one for reproducing the 'exact same' experiment.

This leads to the realization that *one rule for all* in terms of federations is not only unlikely to be achieved but also counterproductive. Different objectives of testing define different objectives for the experimentation required and therefore demand for different approaches to federation. This needs consideration in the future development of federation approaches, both on the technical and also on the governance level.

### B.2.2 Multi-disciplinary aspects

In system design nowadays, industry plays many roles. Vertically, industry tends to think as architects (leaving out the crucial question what constitutes an architect), as technologists, as business designers, or as a business strategist (with the charter to advise or make decisions of crucial investments that will directly influence the viability of technical solutions).

Horizontally, industry tends to think in more or less traditional layers, i.e. as service people, network designers, transport people (pipes) or user-centric researchers that tend to forget or ignore the technical underpinning of the phenomenon to be studied.

Much of this isolation continues in (experimental) evaluation with little involvement beyond the own community. This focuses experimental research to largely component-level testing and experimentation where inbred thinking can easily be afforded. But it also confines thinking towards constructing the *right* facilities to the ones that serve *my* purpose only.

It has been recognized that research towards the Future Internet cannot be confined to isolated technology problems only; that the problem is as large as the usage of the Internet is already today and is bound to become even larger. Hence, we need to make experimental platforms part of a holistic system design process, involving a wide variety of communities. The different objectives of federated

testbeds, while being well argued in isolation towards clearly formulated objectives (see above) must, in the end, serve the problem of solving large-scale issues and questions that surround the Future Internet, allowing industry to stress test components as well as validating new approaches to business, investigate the early impact of new technologies on the behaviour of particular user groups or to study new forms of regulation without requiring the current law to be changed.

### B.2.3 Increasing the pace of evolution

It is the role of research to formulate the visions and to provide the early evidence of them becoming true – it is the role of industry to provide the final evidence in the future marketplace. Industry relies on researchers delivering pre-competitive ideas that can be turned into evolutionary Future Internet products in less than 2 to 3 years time. This timeframe is more compatible with an evolutionary approach.

Nevertheless, both explorative as well as evolutionary questions are required for pushing forward visions that might become commonplace in the future. Whilst the evolution of current technology along well-known business metrics is needed for the reasons above, crucial paradigm shifts must be considered that will question the fundamentals of today's value chain. One such example is that of information-centric networking, which places information at the centre of attention, separating more clearly the concerns of facility and information providers with the potential of changing the way we exchange and pay for data. Understanding such paradigm shifts not only requires fundamental component-level research, but also experimentation at the system-level, supported by a variety of testbed federations, each serving the particular objectives of the tests at hand.

Service and content providers at all levels have to be aware of 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.

### B.2.4 SMEs

Some of today's SMEs will be the major employers and wealth generators of tomorrow. SMEs will benefit by having the opportunity to detect significant trends early, and adapt their development, integration and production activities quickly to the new technologies. Thereafter, they can benefit in the same way as equipment vendors (see above) from the capabilities to trial and evaluate technologies and system solutions in a large scale with real users.

User SMEs who are willing to take the risk of being early adopters will benefit from early access or use of FIRE technologies to accelerate their business.

## B.3 The Research and Academic Perspective

Most of the work in the research community is done in the public domain, often using software that is developed and licensed under open source rules. The reason for this is that in the early stages of developing a novel idea, researchers benefit from the additional inputs that arise from sharing their results and data exceeds the possible loss of proprietary advantage, since there is still a long road to travel before these research ideas become new products, services or businesses. Proprietary advantage still resides in the greater understanding of the practical aspects of the new ideas that the researcher has achieved.

Research experimentation has two different dimensions, which tend to separate the constituency into two parts, which do not always understand each other. The very earliest concepts need to first be proven valid and workable. Experiments with this objective are similar to debugging new software -

reproducibility is key to discovering the cause of unexpected or simply incorrect behaviour. At a slightly later stage, exact testing is less interesting than statistical characterization or even exploring a new application or device in the hands of human users. Here, reproducibility is not possible, and experiment design aims at achieving solid statistics.

Metrics of success for this community are the number of high quality publications cited and the use of the testbeds for their evaluation work.

#### **B.4 FIRE from a Cross Layer / Cross Technology Domain Testbed Perspective**

It can be witnessed that there is a broad notion of testbeds. Very often these testbeds relate to specific technologies, such as a Zigbee sensor testbed, a WiMax testbed, a UMTS testbed, an IMS testbed, an autonomic communications testbed, or a mobile applications testbed, etc. Sometimes also specific (and possibly open source) software toolkits are used to establish globally distributed testbeds, such as the PlanetLab software.

In face of the ongoing convergence of network technologies and application domains, an additional need for cross technology testbeds emerges, i.e. testbeds which cover different network technologies in order to test for advanced handover algorithms and procedures, or triple play and seamless service capabilities, etc. In addition, the time horizon of research projects demands for interworking, migration and evolution experiments, which also cover different technology domains.

In regard to the above, it is important that in addition to the federation of technology-specific testbeds, which is straightforward and in any case needed for large scale experiments, also cross technology domain federation is needed, in order to test, for example, different application approaches on top of converging network technologies, such as P2P versus SIP-based IPTV services. This means that there is the need for generalized resource description, management, control and federation approaches.

#### **B.5 FIRE from a socio-economic perspective**

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

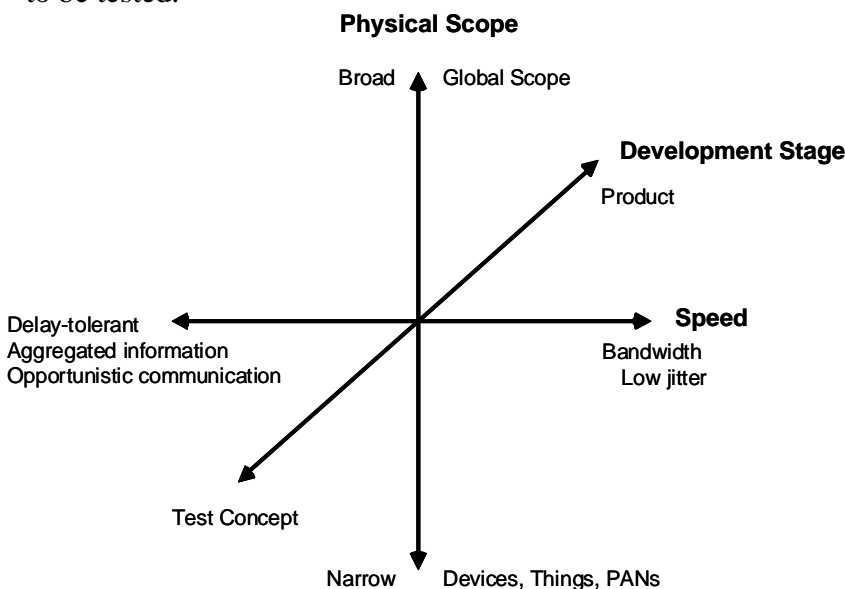
useful public information network aspires to treat all content, sites, and platforms equally. NN is perceived by many as being endangered by policies such as differential quality of service, yet others see NN as a deterrent to innovation and investment in network infrastructure. Opinions are split about 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? What traffic has timing constraints and requires treating with priority?) that may affect NN.

Today, the number of web services publicly accessible through the Internet account for 30.000<sup>3</sup> approximately, this figure is expected to grow exponentially in the Future of Internet scope, motivated, among others, by: the expected grow in net accessible devices and their enclosed appliances acting as service consumers, the improvement on communications infrastructures that will lead to more and better access to the network, as well as, the popularity gained by Web 2.0 concepts of mash-ups open the door for future explosion on business service networks and business flexibility methods and tools. All together, raise important challenges at socio-economical level, such as: the legal context for e-services, the modification of the service consumer and user relation, as well as, trust and security requirements in services and networks interactions.

### B.6 Relative positioning of the testbed projects

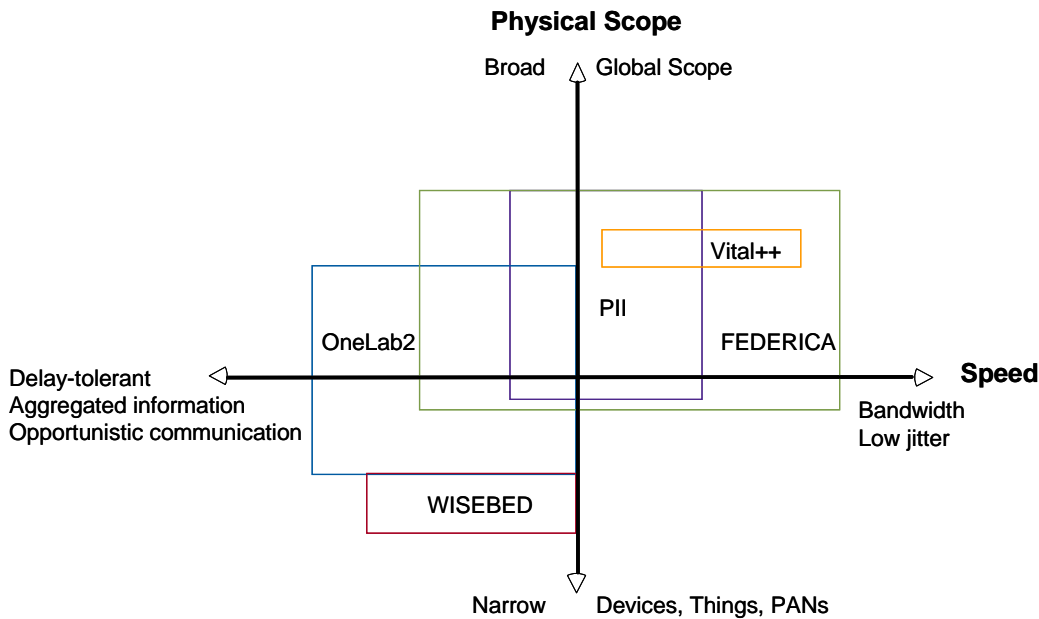
Research and testbeds are clearly complementary. The following portfolio analysis explores how the complementarity works (and what are its limits):

There are many ways to describe the complementarities of the existing FIRE portfolio and to look for additional opportunities to impact the Future Internet. The diagram below shows some dimensions of typical areas addressed by the technology-oriented STREP projects above, which have - to some extent - to be tested.



Considering first the more technical axes of *Physical Scope* and *Speed*, the capabilities offered by the testbed projects can be represented as shown below:

<sup>3</sup> <http://seekda.com/about/web-services>



OneLab2 has global scope through its federation with the three existing (and several future) PlanetLabs (US, Japan and Europe) and at a second level of the hierarchy down to private PlanetLabs and wireless testbeds. It includes the slower, opportunistic and delayed aggregation information technologies, and also research projects in embedded intelligence and wireless technologies (e.g. control facilities that have been developed in the course of the ORBIT wireless testbed).

PII has global scope through its descriptive facilities and registry for the characterization of several heterogeneous testbed facilities provided by European industry. PII provides high bandwidth remote access to these facilities.

FEDERICA is a dedicated Europe-wide network, built using resources from NRENs and GÉANT, thereby enabling global coverage. It forms a platform for Future Internet research, and is therefore a facilitator for other projects to achieve their results. It comprises a combination of computing systems and wired networks capable of virtualization in each element, and will also contribute to the development of research on virtualization. It can host experiments on a wide range of topics.

WISEBED integrates sensor networks wirelessly.

Vital++ addresses the changes in IMS to permit creating overlay distribution networks.

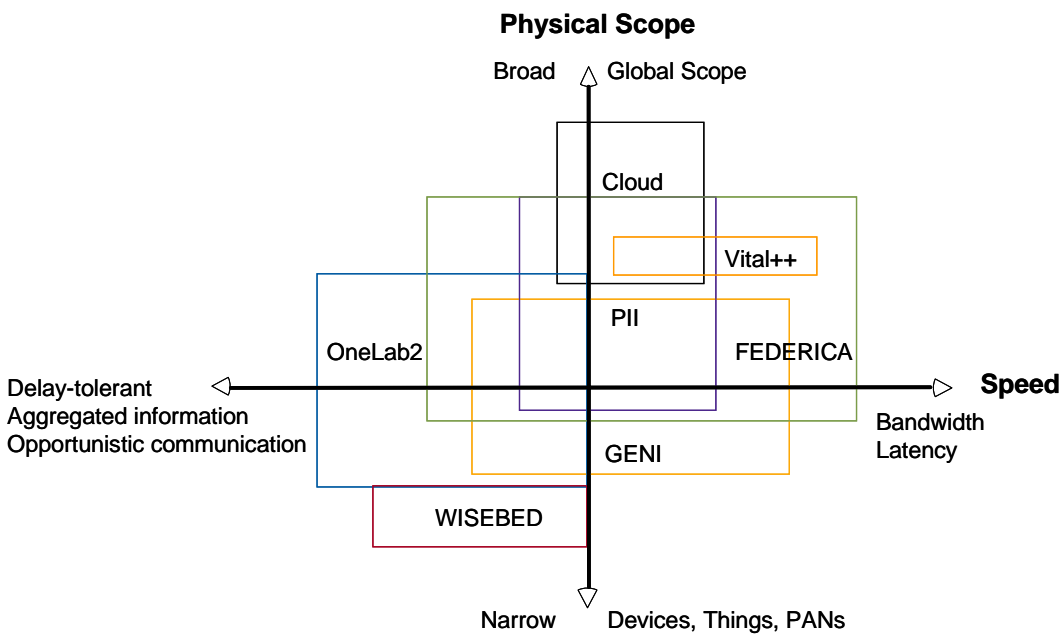
The major center organizing US activities on the Future Internet is the NSF's GENI. Although its efforts are restricted to non-commercial applications at present, it incorporates multiple Control Planes under development. It extends to address both wireless and delay-tolerant testbeds and incorporates several regional high speed optical networks.

Like OneLab2, PII and FEDERICA, GENI provides users with access to infrastructure that is built to support research on the Future Internet. Also, comparing the FIRE testbeds collectively, we can see that they cover both a wider area and a broader set of immediate customers than does GENI alone (see below).

The testbeds can also be categorised in terms of their suitability for experimentation from initial tests of concepts to market-ready solutions. Projects that support rather the initial testing of hardware and distributed application ideas, with a common preference for open source software and public access to results are OneLab2, WISEBED and FEDERICA. These projects have fewer – but still some – concerns about IPR, privacy, security and anonymity of measurements, whereas PII and Vital++ are much closer to market and commercially sensitive. The test brokerage facility provided by PII and

Vital++ therefore includes the extra legal framework needed to provide confidentiality and intellectual property protections for those ideas which can be quickly commercialized by European enterprises and SMEs.

Finally, in the diagram below can be seen the "Cloud" application environment. It represents a layer not present currently in the FIRE facilities, and necessary for FIRE's goal of development of the large scale testing facility for the Future of Internet: the Internet of Services and Internet of Things components, architectures and business models. The Future of Internet has been envisaged as a service platform where providers and users of services, using all kind of devices, consume, produce and supply composed on-the-fly and existing services in a secure transparent seamless manner, enabling both providers and users of to develop new business models by means of the Internet usage. Current projects in the IoS and IoTs are addressing mainly the following areas: (a) service front-ends that enable user-centric service creation and heterogeneous service interfaces; (b) Resource virtualisation to enable service delivery independently of the underlying physical platform; (c) software engineering methods, tools and architectures to tackle the complexity of distributed, heterogeneous and dynamically composed service environments; (d) service frameworks and service delivery platforms to allow services to interact in dynamic service value networks in global service ecosystems. FIRE experimental facilities will have to hold test-beds from these research areas in order to cover the envisaged spectrum for the Future of Internet research Facilities.



### B.7 FIRE and visionary research on the Future Internet

A strong link between experimentation and long-term research is one of the key ideas of FIRE and a distinctive trait with respect to other parts of the ICT work programme. With respect to the research dimension of FIRE projects, this translates into the concept of iterative cycles of research and experimentation, as the way in which FIRE research project should progress towards their final goals.

We stress the fact that FIRE research projects should have a long-term vision, as they must contribute to shape the Future Internet. However, even radically innovative ideas must be tested at early stages by exploiting the FIRE experimental facility that is growing up in parallel. This methodology brings several advantages. On the one hand, early prototyping and testing of innovative, long-term ideas is a premiere way to understand practical implications of the conceived innovations in realistic environments, giving the opportunity of incorporating this feedback in next design cycles. On the other hand, this methodology enables improvements of the FIRE facility as well, either by proposing novel

test cases for the FIRE facility components, or by contributing to extend it through custom components that are required to test the proposed ideas.

Therefore, the ideal target of a FIRE research project should be:

- primarily, to investigate innovative ideas and long-term concepts for the design of the Future Internet, by testing them, since an early design stage, on the FIRE facility;
- to improve the FIRE facility, by identifying Use Cases relevant for testing the proposed research concepts, and/or by contributing to the extension of the FIRE facility.

FIRE has the defining / unique niche of jointly pursuing research and experimentation. Other initiatives either focus on fundamental research and new theories or they follow closely technological and market trends or be primarily concerned with users and applications, whereas FIRE is a bridge between new theories and users/applications.

Theoretical and abstract research has long been disconnected from its real-world application and experimentation. Similarly, experimentation has mostly been seen as an infrastructure issue and as a facility that is good to have but in practice very little and superficially used. Bringing together theoreticians and experimentalists has always been the challenge.

### B.7.1 Cross-layer, multi-level approach

The amount of content and information, and the complexity of the tasks that the Future Internet shall address, require extremely efficient networking solutions and resource management policies. There is a clear trend in the research community to look at including more intelligent features in the core of the network, with respect to the conventional IP model. However, this produces a challenging trade-off to be met, between the performance in terms of throughput provided to users, and the elaborations that the network should perform to accomplish the additional tasks. A possible way to tackle this issue is making the network context-aware, and exploring cross-layer design approaches to share information between the different network layers and functions. Cross-layer design patterns have been successfully applied in the area of mobile ad hoc and sensor networks, for example. Sharing information through cross-layer interactions avoids wasting resources in gathering the same type of information at different layers of the protocol stacks.

More generally, we expect that cross-layer, multi-level approaches will be a key design pattern for the Future Internet design, as a premier way of introducing more intelligent functions at a reasonable price in terms of network operations complexity.

### B.7.2 Content-centric architectures

There is an increasing consensus in the research community on the fact that the Internet is “becoming flat”. Nowadays, users of the network are interested in accessing particular pieces of content, rather than connecting their device with a particular host on the other end of the network. Users generate massive amounts of content that is possibly of interest to other users of the network. Users are actually prosumers, as they both produce and consume content at the same time. Conventional Web technologies as well as P2P overlay networks are used to store, advertise, search and consume content. These solutions (P2P technologies in particular) “blossomed” at the edges of the Internet, as such a design paradigm allowed for quick and flexible development of innovative content-centric networking solutions.

The Internet has been mainly used as a high-performance transport network, while the intelligence required to manage content has been pushed to the edge. Instead, in order to accommodate those user access patterns that already manifest themselves today, and the efficient development of the resulting applications, the Future Internet could provide content management features as a premier service, as

connectivity is provided today. Content-centric architectures can be a key part of Future Internet innovations, offering services such as content replication, storage, search, indexing, to users' application. These services (together with connectivity) should be natively provided by the network, and implemented by enriching the core-network functions with respect to the current Internet architecture. Core devices shall not only be limited to forward traffic at the highest possible speed, but elements should also be identified in the core of the network, which are responsible to implement the functions required to support content-centric services.

### B.7.3 Mobility

This is even of greater importance if one considers the fact that the Future Internet will be largely mobile. The proliferation of mobile devices with rich networking capabilities (e.g. WiFi smartphones) will result in "spontaneous" infrastructure-less networks, formed by mobile users devices, according to the ad hoc and opportunistic networking paradigm. These types of networks will complement the conventional mobile networks of telecom operators, and will provide improved communication opportunities (and, ultimately, reach-ability) to users. On the other hand, a massive penetration of mobile devices with content-rich capabilities will increase the possibility for users to generate content anytime and anywhere. Coupled with the possibility of these devices to be connected to the Internet (although possibly through intermittently-connected DTN access patterns), this will result in an exponential increase of the amount of content to be managed by the network. In addition, any particular piece of content may even be of interest just for a small subset of user, and for a limited amount of time. This means that the content generation, indexing, search and consumption processes should be very dynamic. It is clear that in such a scenario current solutions to manage, index, search and retrieve content (e.g. the Web coupled with Google) will not scale to support such massive amounts of content, and will not provide the level of flexibility and dynamism required to efficiently support the dynamic generation and consumption of content. Therefore, innovative solutions are required that consider mobility, "spontaneous" networking, and dynamic content generation as natural dimensions (and challenges!) in the design of content-centric architectures for the Future Internet.

### B.7.4 Service Research on FIRE facilities

In order to cover all the spectrum of research taking place under Challenge 1, FIRE facilities will have to allow the inclusion of testbeds from research projects of all aspects globally accepted in the Future Internet model: Network of the Future, Networked Media, Internet of Services, Internet of Things and Security, etc. So that, testbeds from these research areas can be federated within the FIRE facilities to consolidate results and validate their research on Future Internet.

Focussing on the IoS and IoT R&D projects, in order to allow the necessary large scale and repeatable experimentation of service-based models and components, it is recommended the extension of FIRE facilities geared towards services so that current software and services projects can advance through controlled experimentation and optimization testing.

### B.7.5 Resiliency

Resiliency is an important challenge for an infrastructure combining networking and computing/storage infrastructures. End users will be faced more and more to the unavailability of services due to the increasing complexity of service based applications, single point of failures, etc. This will generate a lot of frustration for the end users, which should not be underestimated.

### B.7.6 Social networking

Social networking is a strong trend that needs to be reflected also in the FIRE strategy. Not only should it be studied and incorporated in the complete picture of Internet research, but also understanding its mechanisms can lead to new approaches in terms of layering that can benefit even disruptive solutions.

### B.7.7 Socio-economic aspects to be strengthened

FIRE should be the playground for interdisciplinary testing including socio-economic aspects. Far too often the socio-economic aspect in ICT projects is restricted to the research carried out in the beginning of the project in order to create understanding of the state-of-the-art. It is a challenge to bring together people from different disciplines in a meaningful and effectively resulting manner, but further efforts towards this direction should be stressed. The more complete and deeper understanding of the Internet(s) and its (their) multiple missions need to be developed and shared. Means for strengthening the socio-economic aspect should be fostered.

### B.7.8 Technological issues (non-exhaustive list):

- Autonomic communication for content distribution
- Autonomic governance; dynamic capability of exchanging information
- Context awareness
- Data confidentiality
- Network solutions, based on programmable components
- Resilient, future-proof communication software
- Robustness and safety, to adapt to changing environments, ...
- Scalability issues
- Self-\* properties (including self-configuration of network elements)
- Semantics
- Software autonomy
- Monitoring and measuring

## B.8 The FIRE facility

The FIRE experimental facility will be built by gradually federating existing and new testbeds<sup>4</sup> 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.

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. This federation, however, is created with a set of clearly identified objectives in mind, making the federation useful to begin with. Such objectives can be achieving scale, access to heterogeneous resources, open exchange of data, testing of new pre-

---

<sup>4</sup> A testbed here is to be understood as an environment allowing rigorous, transparent and replicable testing and experimentation in research and development projects

commercial products, interoperability testing of hardware and software component. It is important to note that not all of these objectives can be realized in a single federation since they can be contradicting. Hence, there can be a need for several federation concepts, depending on their set of objectives (and the constraints in achieving them).

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.

Some key challenges for the FIRE facility are:

- **Federation:** including access to the facility, interoperability framework and Control Plane. A FIRE facility builder's task is to develop mechanisms for federating testbeds and to contribute to the FIRE facility as a whole, by defining interfaces to different stakeholders. Research projects having an objective view on these interconnecting mechanisms and metrics of testbeds are encouraged to contribute to the best practices and neutral processes for building the FIRE facility.

There will quickly be situations where multiple projects wish to make use of similar facilities but have different approaches to implementation. Further, we may wish to enable others outside FIRE to build versions of functionality available within the facility - time and time again this model has proven an accelerator in Internet/Web development. In the FIRE facility's case, the interconnection specifications need not be heavyweight or numerous, however they may provide a valuable "glue" between projects and could subsequently be pushed to standards bodies such as the IETF, ISO or W3C more easily than ad-hoc documents.

- **Governance of the federation:** The governance of the federation is critical to the success of FIRE (short, medium and long term). There are many models which could be adopted; however an inter-project (or simply independent) governing committee for the FIRE facility is probably needed, which includes (directly or as subcommittees and rotating on some democratic principle) at least:

- overall strategy / direction
- technical architecture / specifications
- operations

Such a coordinating committee could provide a way to push results between different projects and out to general usage. Organizations such as those running GEANT and EGEE would be good models for the operations elements.

- **Incentivising 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.

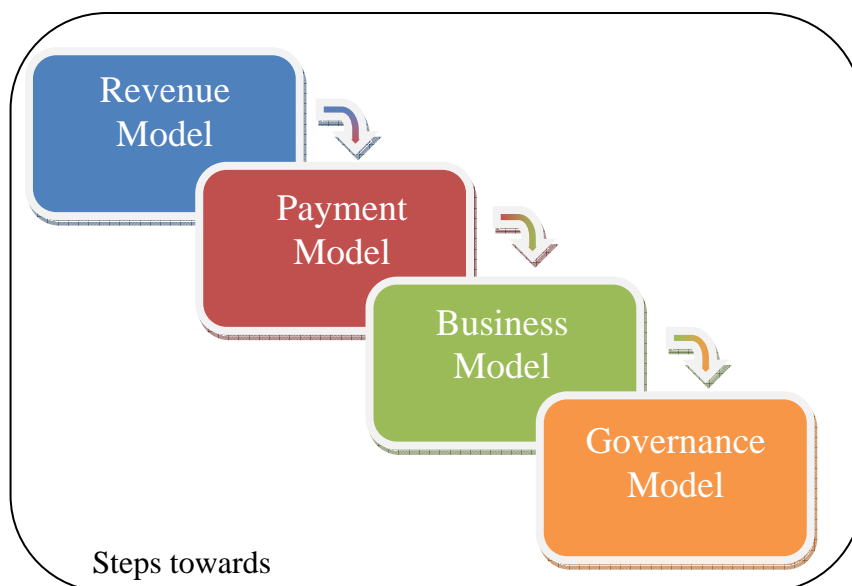
A transparent classification and rating of the testing sites would improve for everyone the overall quality of the offered services. By bringing the individual test sites on the same line, a competitive situation is induced between the sites. This alone is an incentive for the sites to stay up-to-date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

- **Sustaining the federation:** 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.

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 GÉANT. In order for the experimental facility to become self-financed by the stakeholders in the medium term, in the start up phase support is needed for:

- 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;
  - managing and operating the FIRE experimental facility;
  - provisioning some dedicated resources for the core of the infrastructure.
- **Business models towards self-sustainability:** The sustainability of the FIRE facility is in the interest of all the stakeholders. It ensures that the investment in FIRE is secured and that the experiments can be repeated over a large length of time. Ideally the facility will be self-sustainable, that is to say that it does not require constant cash injection from the European Commission or national governments to remain operational. This is obviously beneficial in that this money can be invested elsewhere. In addition it incentivizes efficient management and a user-focused approach.



- **Sustaining the measurement data:** the sustainability of the measurement data and its metadata is also an important issue. The huge amounts of information, their heterogeneity and the expected privacy concerns are difficult issues, and solutions for their management (within or outside the FIRE activities) should be found to set up a common data repository. It is also necessary to know how the results were obtained; i.e under which experimental conditions and testbed configurations.
- **Virtualization:** running concurrent experiments, whilst ensuring no interference between slices, and no breaches of data security between experiments running on different slices.

- **Monitoring:** collecting data and making it available. Topics include:
  - archiving and publishing fairly detailed data from experiments, but respecting privacy (by anonymising the results) where necessary
  - providing the capability to allow others to modify (for example) algorithms and repeat the experiments, without having to build the hardware themselves.
- **Reproducibility:** To the extent that (in the absence of very tightly controlled test and background conditions) it is at all possible for any repeated experiment on a large-scale to give exactly the same result, FIRE should seek to promote methods for supporting reproducible experimentation. However, in many cases what is needed is to be able to repeat many times a test scenario in an automated way, without necessarily expecting to obtain the same result, since changes will have been made to hardware, software and configurations in the meantime.

For example, when experimenting with a new service offered as an overlay service, the tests will be designed to study the characteristics of the new service under a wide variety of demands and conflicting resource constraints, possibly with uncontrolled real-world inputs, and over a long duration. Repeatability is not possible in these cases.

The end results will generally be summarized and published in papers that explain the statistical significance of the measurements and the scientific quality of the test will be the reproducibility of the basic claim, not exact repeatability.

On the other hand, for Use Cases that involve prototypes which are almost ready for market launch, but require performance testing, detailed user testing, or interoperability testing to eliminate potentially fatal snags, then repeatability is a valid objective. However, these tests are usually short.

- **Benchmarking:** assessing the results.
- **Security:** robust and secure facility.
- **Economics** of the facility, for the users, for the operators; conditions of use. The cost of the physical access to some FIRE testbeds can be a problem. For example, an obvious way to gain access to FEDERICA from OneLab2 nodes would be via the NREN networks. However, it is not feasible at current pricing rates to finance experiments, where 10-100 Gbit/s traffic (each) from 50-100 different network locations would like to reach the GÉANT core simultaneously. Therefore the new FIRE facility should directly couple these infrastructures so that access-to-core traffic experiments should not generate extra costs. These experiments should be an inherent part of the single facility.
- **Legal** responsibilities and liabilities, IPR, etc.
- **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.

### B.8.1 Next steps

Since new facilities will soon need to be integrated into the overall FIRE experimental facility, the existing core facilities must achieve consensus quickly on (i) how they can be federated / integrated / interconnected / loosely coupled, (ii) how new facilities can be integrated, and (iii) how their resources, services, interfaces, access capabilities, etc. can be advertised to users, and each other.

As a general rule, the core facilities that have been established should continue to be supported and extended through bringing new innovative testbeds into the federation and/or adding new interfaces, capabilities, paradigms, etc. It would not be cost effective to start again from scratch.

## **B.9 The strong relation between FIRE Research and the FIRE Facility**

FIRE simultaneously (i) investigates new and possibly unexpected Internet concepts; and (ii) gradually federates existing and future testbeds to become an integrated, sustainable, dynamic, large scale experimental facility on Future Internet technologies.

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 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.

## **B.10 Encourage the re-use of already-developed software/protocol stacks**

Incentives for utilising past projects' outcomes, such as software, protocols, platforms, etc. should be introduced in FIRE. So far, projects proposing new approaches and starting from the scratch have tended to benefit in the evaluation process over extensions of so-called "legacy" projects. It would be coherent to support the long-term vision by building on the existing results, especially in the case of FIRE.

## **B.11 Additional FIRE Use Cases and scenarios**

The use cases described in this Section are just examples, to give a "flavour" of the future FIRE activities. They should by no means be seen as an exhaustive compilation of possibilities. More potential use cases are summarised in the annex.

### **B.11.1 Cloud Computing (using OneLab2)**

Different business models have emerged in CC that have parallels for increasing the utility of PlanetLab and FIRE's testbeds more generally:

- hardware as a service, in which you purchase what you need for as long as you need in standard increments of storage, cpu power, and duration, as long as you are can accept the supported Oses and system standards.
- software as a service, same terms and conditions, but without even caring about the nature of the underlying hardware you can lease large database tools, query systems, standard transaction processing environments.
- service offerings emerging directly from the above two, in the hands of small companies building solutions with no hardware and minimal software investment (e.g. Salesforce

automation). This can be the fastest growing source of competition for European companies like SAP.

From the point of view of those pursuing distributed computing, today's "clouds" are just the same old client server model on steroids, with the extra business advantage that at present hard disk storage economies don't support high I/O bandwidth, leaving a system designed for transaction performance with lots of spare storage to sell to less-demanding uses. However they are all based on proprietary standards. The Amazon cloud, the Google cloud, and the Microsoft cloud will have to develop common interfaces, and the applications world will have to develop ways of using them more interchangeably.

FIRE can play a key role in moving computation out of the fortified, proprietary clouds and into a smoother range of interoperable distributed services that are building blocks of applications that go beyond today's cloud successes.

Clouds today have no special tools for reaching out the edge of the network (cellphones and laptops). Clouds today use ultra fast networks within their complexes, but assume least common denominator networking to the outside, or perhaps take advantage of an Akamai (fairly rigid, and high-end) solution to global content distribution. Connecting P2P content distribution to a high performance source to give more interactive behaviour will need to be tested, and FIRE can help.

Making this kind of high end application more aware of external network capabilities in real time can be facilitated by the kind of tools that the OneLab/Federica is now deploying. Middleware to connect the cloud to the lowest layer dynamically is the subject of 3DNETS, a Call 4 proposal that will explore federation of FEDERICA and ONELAB facilities (participants come from both projects).

### B.11.2 Cloud Computing (federating OneLab2 and Grid 5000)

**Background:** Service Oriented Architectures (SOA) will take a predominant place in the future of the Internet. However, deploying and managing services requires large-scale data centers to host them such as those developed by Amazon, Google and others. These large-scale data centers are the cornerstones of Cloud Computing, aiming at providing the instant delivery of resources to host new services. This real-time aspect is the essence of the Cloud Computing paradigm, especially when comparing to previous ones such as Grid for which resources are usually provided by batch systems. Interactions between a very large number of clients (usually PCs spread over the Internet) interacting with Cloud systems (data centers) will have to be studied, thus stimulating new research into distributed processing. Operating System virtualization will provide the flexibility to design new Cloud architectures that will be scalable, self-managed and less power-hungry while keeping the instant delivery within the limits specified by service level agreements. This is indeed a research challenge per se since today most of the Cloud systems provide a best effort policy to customers. This is largely due to the lack of tools, methodologies and algorithms to map SLAs into reliable Cloud architectures. Research into Cloud Computing will need large-scale testbeds.

Within the FIRE initiative, and its supported testbeds, it is today feasible to allocate a set of independent geographically dispersed PCs. However, accessing large scale servers having both computing and storage resources, in order to experiment Cloud systems under realistic conditions, is not yet possible.

**Use Case:** Synergy between OneLab2 and Grid'5000 (see sub-section 7.1.1) leads to joint research activities in the area of Cloud computing. Grid'5000 provides today and for the next four years, a federation of at least 9 large PC clusters having in total 1483 machines with almost 5000 cores. The full reconfigurability of Grid'5000 will allow experiments of Cloud systems without imposing any particular technologies. However, this is the theory. In practice, allocating two different testbeds (OneLab / Grid'5000) with their own administrative/security policies and network configurations, still

needs a lot of technical works and skills. A user of such a federated testbed should simply specify its needs: “Next Tuesday, I would like to have 1000 PCs geographically dispersed interacting with two interoperable Cloud systems having each 2000 cores running the Eucalyptus and Nimbus software stacks hosting some specific services”. Today, this is not feasible although there is a clear need to carry out such experiment on a daily basis.

### B.11.3 Emulating AS Relationships (using FEDERICA)

**Background:** Data in today’s Internet is transported through an interconnected web of ‘Autonomous Systems’ (AS), which are usually provided by some facility provider (e.g. in the core) or content provider (at the edge). Transporting data against a defined cost model is called ‘transit’ while exchanging data in a swap type fashion is described as ‘peering’, i.e. the data exchanged is not paid for. The current transit and peering relations in the large-scale Internet have evolved as a careful set of business relations, reflecting the current tiered value chain of core and access network providers.

The establishment of today’s relationships, however, are reflecting the current pre-dominant usage models for the Internet. Large content providers on the one hand provide vast amounts of data to many providers who provide the required clients that retrieve the data (e.g. end users retrieving YouTube videos).

Changing the usage models, however, challenges the way these relationship are defined. For instance, the observed locality of certain content leads to increased peering among same-level providers, avoiding costly transit charges to high-level tier providers. It can be expected that many new usage contexts of future services, such as sensors, context-aware services, localized streaming and many more, will continue to challenge these relationships. But also technical advances and proposed new solutions for the Future Internet challenge the relationship. For instance, many data-oriented network approaches make use of locality through increased caching. This is likely to have an impact on the AS relationships which will in turn influence the viability of the approaches themselves.

**Use Case:** A testbed provides the controlled (enough) environment to experiment with various AS relationships. Hence, this Use Case envisions the programmatic setting of ANY AS relationships, while conducting experiments of various service and/or networking technology. The testbed facility should not only provide the technical means to ‘reprogram’ the AS relationships but also to visualize the resulting AS topology and give a basic economic understanding of the new relationships.

Moreover, a formal policy language would be provided that can allow for such ‘reprogramming’ on a high enough abstract level rather than dealing with (for example) BGP-like policies.

There are many specific Use Cases that would make use of this programmable structure. Examples of these are:

- testing new inter-domain publish-subscribe solutions in the presence of different AS relations, e.g. efficiency of inter-domain rendezvous solutions
- testing Content Delivery Networks solutions in the presence of changing AS relations, e.g. along a well-developed scenario of industry transition
- testing federated sensor solutions in the light of localized caching (the extreme of localized traffic)
- testing Cloud Computing scenarios under changed AS relations

Technically, the proposed Use Case calls for a virtual topology that does not rely on the Internet as we know it. A first iteration, however, would build this virtual topology over the existing Internet, virtualizing the inter-AS relations into an own set of relations under full control. A second iteration would set these virtual ASs onto controlled link layer resources, wherever possible, e.g. through dedicated networks such as provided by FEDERICA.

#### B.11.4 Public safety

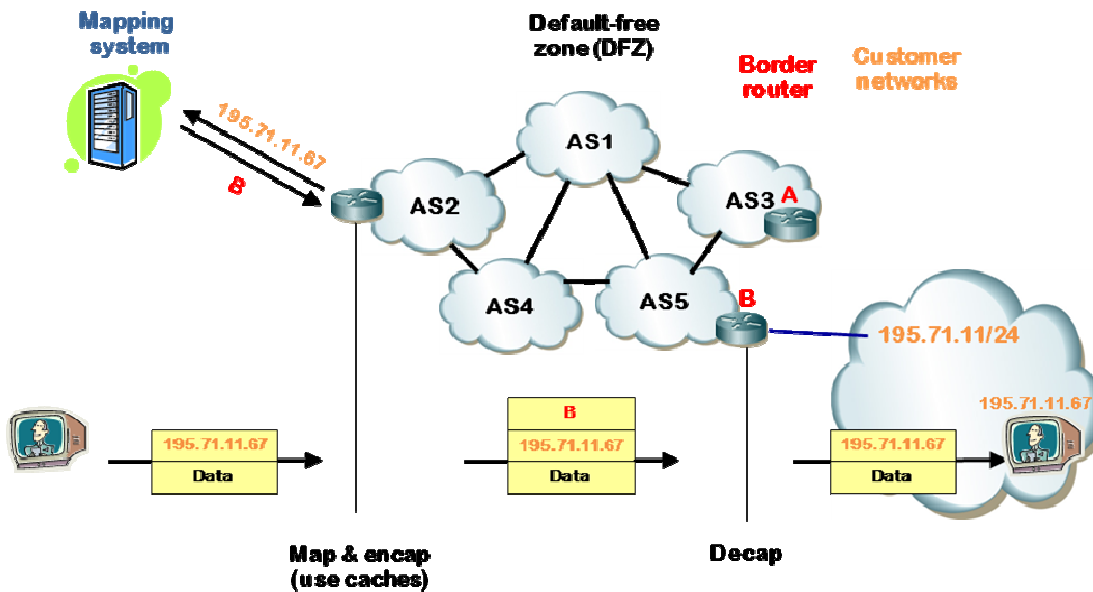
**Background:** Public safety requirements and user scenarios are considerably different from the requirements and scenarios of public communication and information systems, especially in terms of resilience, interoperability, manageability and security. European public safety is currently using a number of different communication technologies such as, TETRA, and TETRAPOL. Moreover, new technologies for public safety such as TETRA-TEDS, and WiMax are in different stages of active development processes. Disaster relief operations, requiring the involvement and coordination between multiple agencies, are particularly affected by the lack of seamless interoperability, or easy control of large-scale and heterogeneous PSC networks. Thus, the emerging concept of network centric public safety operations, the foreseen introduction in the future of new communications technologies for public safety, and the development of heterogeneous networks for public safety call for testing environments and capabilities that could be only provided by international federation of resources, such as federated test beds established through the FIRE initiative.

**Use Case:** In our vision of PSC, all types of networks available in a disaster relief scenario will be used opportunistically, and will complement proprietary networks of emergency operators. Coordination and integration of the different types of networks call for large-scale heterogeneous testing environments. One important requirement for a PSC testbed is to provide a range of different communication technologies, whose availability can be controlled dynamically during the course of the test runs. This allows us to test the integration of such communication technologies in very dynamic environments, which is typical of PSC scenarios. Another key requirement for a PSC testbed is large scale, to realistically test communication solutions in relevant physical regions and to scale up them significantly with the number of actors involved. It is also very important that the testbed nodes be configurable and controllable, and that virtual users can be defined in the tests. This allows us to differentiate between the different types of actors involved in a PSC scenario (rescuers, citizens, etc), who clearly have different requirements in terms of communication, and who must be granted different priorities and class of services by the tested PSC solutions. The heterogeneity, complexity, dynamism and configurability requirements of PSC tests call for a testbeds whose complexity is typically beyond the effort of individual projects. Therefore, the growing FIRE facility is a natural candidate to host this class of experimental research efforts.

#### B.11.5 Mapping Systems for Future Internet Routing (using G-Lab)

**Background:** Current Internet routing does not scale since the number of entries in BGP routing tables steadily increase.

**Use Case:** The core idea to overcome this problem is hierarchical routing in the sense that the core of the Internet is only required to forward packets with core addresses. If customer networks want to send packets to a destination in other customer networks, ingress nodes to the core Internet need to consult a mapping system for a specific destination node (195.71.11.67 in the figure), the corresponding egress node (B in the figure) and add this address to the packet, e.g. by encapsulation. This technique is also known as map & encaps. The packet is forwarded to the egress node and de-capsulated so that the original data packet arrives at the destination. There are many proposals for Future Internet routing and this principle, also known as the locator/identifier split (Loc/ID), is common to most of them. Therefore, one objective of G-Lab is to design such a mapping system in a way that it is robust and requires as little support as possible from central authorities. In addition, security, consistency, and performance are key requirements for the desired mechanism. We intend to develop ideas, simulate them, and test them in the experimental facility.



### B.11.6 Locality of P2P traffic - University of Würzburg (using G-Lab)

**Background:** The large amount of data exchanged in overlay applications is a significant source of costs for Internet Service Providers (ISPs). Overlays typically span the networks of several ISPs, and due to the logical separation of the overlay and the physical network topology, content is often exchanged between end users that reside in different ISPs. Such inter-domain traffic leads to interconnection costs for the ISPs. Traditional traffic engineering methods are less sufficient to satisfy simultaneously the provider’s goal to maximize network usage and the user’s goal to maximize his utility, since most of current overlay networks are deployed obliviously of the underlying physical network.

Within this context, the FP7 SmoothIT project (<http://www.smoothit.org>) proposes a new traffic management mechanism termed Economic Traffic Management (ETM), which provides for an incentive-compatibility in those interactions foreseen between overlay applications, the underlying ISP networks and the end user, in order to gain the following measurable impacts: (i) Cost saving for ISPs, through lower operation costs (due to ETM-based traffic engineering), lower interconnection costs (since traffic can be kept inside an ISP’s domain), and lower capacity extension cost (since capacity requirements can be forecasted with much higher accuracy), and (ii) Better Quality-of-Service (QoS) for overlay-based applications across ISP domains, due to the usage of ETM-based traffic engineering. This leads to an improved media consumption experience (i.e. Quality of Experience - QoE) for end users.

**Use Case:** Within the SmoothIT project, a number of ETM mechanisms are developed and evaluated by means of simulation and analysis. In addition, the ETM mechanism is implemented in a testbed considering the underlying SmoothIT architecture and protocol concept to exchange information between overlay and underlay networks. The aim of this Use Case is to provide a proof-of-concept and validate the results and the performance gain due to ETM in a testbed. This will probably happen in addition to the SmoothIT-planned internal and external trials. In particular, the reduction of costs due to P2P traffic in inter-domain environments as well as the impact on the application’s performance is quantified.

In this context, a worldwide federated testbed is beneficial for various reasons. First, it allows for testing large-scale overlay networks which helps to validate the results for realistic scenarios and to generate realistic traffic pattern within the overlay. Second, a federated testbed allows for including

various ISPs and various underlying physical substrates which will demonstrate the ETM concept itself across domain boundaries in support of co-operating and competing providers in an interconnected heterogeneous network environment. In particular, it is intended to federate the SmoothIT trial network with the G-Lab testbed.

#### B.11.7 Congestion aware routing and network navigation – A Use Case from the Computer Research Institute of Hungarian Academy of Sciences

**Background:** Internet traffic is increasing in an unprecedented pace. According to Cisco Research, it will increase annually 46% in the period 2007-2012. This networking trend is driven largely by the increasing use of video and Web 2.0 social networking and collaboration applications. The combination of these technologies represents what is known as visual networking. Video on demand, IPTV, peer-to-peer (P2P) video, and Internet video are forecast to account for nearly 90 percent of all consumer IP traffic in 2012. This traffic is then not originated from large corporate server farms well connected to the backbone; it is generated in the access network. The server-client approach is no longer applicable and users need symmetric access to the network, rather than asymmetric like in current A-DSL. This involves not just an upgrade of the upload speeds of users and deployment of fibre to the home; it is also going to change how routers handle the traffic globally over the network. The access network will lose its tree like topological structure and becomes a complex network with a large number of loops and shortcuts like our current city street maps. Instead of minimizing hop count, new routing protocols should navigate packets just like current traffic information systems help drivers to avoid congested street segments.

**Use Case:** The Use Case exploits federation between FEDERICA and OneLab2, in order to test various routing strategies in the FEDERICA's virtualized environment. Strategies can involve new routing principles reacting to congestion on links and also strategies, where end users can select specific paths for the delivery of their data packets. OneLab2 is ideal for the testing of traffic generation and packet navigation strategies.

**Realization:** The current PlanetLab/OneLab infrastructure supports the testing of various traffic sources in the access network. However, it explicitly prohibits the transfer of large amounts of data coming from high throughput sources between the nodes. The new FIRE facility should allow experiments where multimedia traffic from the access part (OneLab2 nodes) can flood the core network. The FEDERICA infrastructure supports the testing of complex routing strategies needed to navigate a flood of multimedia traffic through the core network and could therefore provide an ideal testing ground for various routing concepts.

The physical access to FEDERICA still has to be resolved. An access from OneLab2 nodes via the NREN networks is considered the best technical solution, but it is not feasible at current pricing principles to finance experiments, where 10-100 Gbit/s traffic (each) from 50-100 different network locations would like to reach the GÉANT core simultaneously.

**Commercial interest:** This use case is important for IPTV providers, P2P video service providers, fibre to the home network providers, intelligent home applications and large scale Internet of Things (RFID) players.

#### B.11.8 SOA Scenario for Engagement with FIRE Testbeds

**Background:** The ideas behind distributed computing have been evolving continuously for the past 30 years. In each decade it seems a new distributed computing paradigm is invented which will deliver the holy grail of pervasive, secure, high performance application solutions for the public and private sectors. The latest paradigm was entitled *Grid computing*. Over the decade, the ideas behind Grid

computing have gradually broadened out and been renamed Service Orientation. It is important to note that while the Grid very quickly became associated with particular middlewares, Service Oriented Architectures (SOAs) follow much more a distributed application design methodology than a prescriptive set of technologies.

Many different communities became involved in Grid computing. With the exception of the Particle Physics community (who, thanks to the very large data transfers which characterise their computing needs, experimented in detail with bandwidth reservation and high performance data transfer technologies), very little engagement with the networking community and those defining the underlying networking protocols has occurred. Most SOA applications today use the network as is (generally with IPv4 and no parameter tuning).

It is clear that the number of large-scale applications built by both the public and private sectors will increase markedly over the coming years. Europe has led the way in the use of Web Services to build SOAs and much of the groundwork for the creation of complex SOAs built from large numbers of cooperating, widely dispersed services has been undertaken in a number of Framework 6 and 7 projects.

It has been clear for some time to developers working on the higher level service technologies employed today that the underlying network protocols greatly influence the performance, security and scalability of their applications. In general, these developers work around network issues and have little or no engagement with network researchers.

**Use Case:** We therefore propose that a key FIRE scenario should be a collaboration between network researchers, service technology developers and service-based application developers to explore how next generation networks can support technology developments. A series of large-scale representative Use Cases from the business and scientific sectors should be selected and used to explore next generation network research ideas. In this context PII in particular can be a key resource along with the test facilities available at the NREN level. Example applications may include: large scale data access and integration for a multinational business wishing to join disparate and widely distributed data sets together for business advantage or the processing of scientific datasets requiring data and information from widely distributed data services, as is common in the bioinformatics sector.

This will only succeed if all three communities work hand-in-hand. It is almost certainly the case that those working on SOA technologies will not know what questions to ask of the networking researchers and vice versa. However, by bringing these communities together, Europe can harness our disparate skills to drive the Internet forward to our mutual benefit.

#### B.11.9 Media content delivery service

**Background:** The current media applications are significantly affected by the limitations of the network technologies which prevent the efficient delivery of the high quality media content to end users. Media formats, resolutions and compression algorithms used nowadays to provide the content are unable to guarantee the best possible performance and satisfy the users' needs. A tremendous amount of work has been carried out to increase the Quality of Experience and, as a result, a set of innovative solutions have been developed. However, some of these solutions (e.g. 4k/8k video) put highly demanding requirements on underlying transport network.

The network plays a key role in all aspects of content distribution. This has a special meaning especially in the context of media streaming where the network parameters like latency, jitter or achievable bandwidth are at utmost importance. In the past, the Internet was mainly used for simple information exchange. With the advent of new media technologies and increased users' engagement in the content creation and delivery, the Internet has evolved towards the virtual environment for media and communication services. Media became an important building block of the current Internet.

**Use case:** Research on New Media paradigms requires an analysis of all relevant network constraints and may influence future network architectures. The testbeds should allow the validation and testing of newly developed media technologies. In particular, this includes:

- new media applications,
- content distribution, and
- media and network convergence technologies.

The new media applications, like 8K, digital cinema, multi viewpoint imaging or stereoscopy must be deployed and validated in “feel-real” environments. This allows the researchers to optimize coding, compression and adaptation algorithms to guarantee the best possible Quality of Experience.

The need for the distribution and broadcast of high bit-rate media content stimulates the development of new network technologies. Due to the greedy nature of high resolution media streams the technologies must provide optimal use of available networking infrastructure, especially optical networks. The research on this topic should take into account the state-of-the-art solutions available nowadays, like centralized and federated approaches (content owner controls the content distribution) as well advanced peer-to-peer networking (p2p, p4p, etc.). These solutions should be also scalable in terms of resolution, e.g. starting from low resolution images for mobile terminals to ultra high resolutions. The test-bed should allow to deploy the above mentioned systems and validate them taking into account the backbone and access performance.

The convergence of network and media layers will bring significant benefits in terms of the management and maintenance of large hybrid infrastructures. This new layer will allow easy service discovery and allocation of both, media and network resources. It should be considered how the extensions may be applied to the existing network protocols to carry media information to support seamless single-step allocation of networked media resources.

The rest of the Use Cases are not described in the same level of detail as the previous ones, their relevance to large-scale experimentation over federated testbeds has to be more elaborated upon.

#### B.11.10 Adaptive monitoring - Fokus (using OneLab2)

The ANA project develops - as part of its Autonomic Network Architecture - a monitoring infrastructure for self-managing autonomic networks. Such monitoring has to be adaptive and programmable with measurement functions placed and configured dynamically in the network. A set of monitoring functional components is dynamic configurable and with self\*-properties. The usage of OneLab2 testbed allows for validation and practical experimentation in real-world like environment. A testbed with nodes based on the ANAcore is planned for PlanetLab Europe. Those nodes allow the experimentation with the functional composition concept.

#### B.11.11 Congestion control - Fokus (using OneLab2)

The spontaneous phase synchronisation of pulse-coupled oscillators is a well known phenomenon in biology and physics. An application of this synchronisation property of oscillators to networking problems is of interest for the following reasons:

- The observed synchronisation property is based on emergent behaviour. No configuration or management is required.
- Once archived, synchronicity corresponds to a stable equilibrium. Small variations of oscillators do not dramatically change the behaviour of the whole group.

Recent research in the networking area has investigated in the question if pulse coupled oscillators can be used e.g. for time synchronisation in Ad-hoc networks. With the planned set of experiments Fokus targets new and different applications of the synchronisation property. The question behind the performed experiment is if pulse coupled oscillators can be applied to congestion control in IP based networks.

#### B.11.12 Outsourcing of testbed environments - (using PII tools, including Teagle)

An organisation (for example an SME in the ICT domain) wants to develop and offer a new service. The customer uses the PII services to browse for available functionalities within the federated FIRE testbeds. It formulates a request for a testing facility that meets its needs in terms of required interfaces and platform capabilities. The PII tools provide a proposal for where, when and how to support the development and testing of the new service.

#### B.11.13 Test of services and applications - (using PII tools)

An SME or industry in the domain of ICT has developed a geographically localized information service, including the user interface on a mobile device. This service has been implemented by a service provider and an operator (which can be the same company) and is successfully deployed in a European country. One of the stakeholders (SME, service provider or operator) wants to evaluate the potential extension of the service to new markets (new types of information and new countries), by a live experience with different user groups. The PII tools and services, and the PII federated testbed, are well suited to identify and recruit the necessary user groups and then to setup and execute concurrent tests at the same time period in different locations. The results of the tests are collected and analyzed to bring out the behaviour and expectation of the final users, to define the necessary adaptation of the service characteristics to the country, and to estimate the benefit of a new deployment with technological data.

#### B.11.14 Interoperability - (using PII)

Interoperability here relates to equipments implementing different specifications and requiring adaptation means and tools for their interworking; in particular between fixed and mobile network domains, including:

- Performance test of new services, networks, network components, terminals over different technology networks.
- Tests of specific functions across different networks, such as:
  - QoS allocation algorithms
  - Terminal/networks protocols
  - Network procedures (handover, roaming, addressing schemes, unlicensed access)
- Set up of different technology access networks to host the experiments

#### B.11.15 Re-use of test suites - (using PII)

The reuse of test suites contributes to achieving testing efficiency as it helps to avoid duplicating the same work, each time new tests are about to be hosted on the PII platform. The wider the test suites portfolio, the higher is the possibility of finding a suitable test suite. Therefore, PII intends to develop a

broad range of test suites that will remain reusable by its stakeholders as reciprocal benefit for being member of the testbed federation.

#### B.11.16 Certification - (using PII)

Certification is the process of monitoring the operation of a system against a set of functions that are usually specified by one or several standardisation bodies in a number of technical specification documents. PII will not develop its own standards but will follow existing ones, while the user request should comply to the performance layering that is in force by the related regulatory body for the system under test, otherwise the test request must be rejected. Certification is an important procedure through which a number of vital aspects concerning integration of the system under test within networks can be monitored and quantified. Such aspects include:

- **Interoperability of networks:** A system is checked for compliance to a minimum set of functions, which if satisfied guarantee a certain level of interoperability of the system with neighbouring components in the network.
- **Compliance to standards:** A system is certified against support of a set of standards, usually expressed as protocols or protocol stacks. As new procedures are added to networks, new versions of protocols/protocol stacks are specified. Certification is therefore important for testing a system's compliance to the new protocol features.
- **Compliance to network procedures:** The advent of mobile and wireless networks has created the need for approving networks and individual network components for the support of procedures such as roaming, handover, mobility management, support of location based services. As in the previous cases, certification is applied as a set of test suites, which can be executed in order to categorise the system in several compliance levels with regard to the procedure under test.
- **Compliance to a number of performance issues:** Performance of networks or systems is defined as the capability of the network/system to sustain normal operation under certain stressful traffic conditions. For example, the network may sustain without losses a number of users, exploiting in parallel communication services or applications.

Certification is important for guaranteeing best user experience. Especially bearing in mind the heterogeneity of Next Generation Networks technology concerning physical communication interfaces, protocols and services, certification is an issue of vital importance that guarantees networks and services interoperability. It is therefore anticipated that the certification process of the PII federation will be useful not only for operators and network manufacturers but mainly for SMEs, which cannot afford setting up experimental testbeds to host certification experiments. Examples of systems requiring certification are:

- Test equipments implementing a number of testing capabilities, such as traffic monitoring, simulation, emulation or even load test need to be certificated prior to be launched on the market
- Network installations developed by the industry need certification against the number of parallel users they support, interoperability of protocols and procedures with external networks and capability in hosting certain services
- Manufacturers of network components are interested in testing the performance of the systems they produce as well as to guarantee a certain level of interoperability

### B.11.17 Service Oriented Architecture validation - (using SOA4All)

Background: SOA4All provide a comprehensive framework and infrastructure that integrates complementary and evolutionary technical advances (i.e. SOA, context management, Web principles, Web 2.0 and Semantic Web) into a coherent and domain independent service delivery platform. The main objectives of the project are to introduce the following key enablers to realize a web of services interconnecting billions of services:

- Scalability to master the very large and meeting the challenge of dealing with billions of services.
- Ease of use to allow end users not only to interact with services but also to create services.
- Discovery to find the “right” service among the billions of services offered.
- Dynamic composition of services to create the networks of services underpinning business networks and aggregating seamless solutions by orchestrating a network of services.
- Increased robustness to respond to change. Services will appear, disappear, change location, adjust their usage model from free-use to pay-per-invocation, will be blocked, out of service, be inspected and commented upon – with dependent services being able to act quickly to changes.
- Context-Awareness to support collaboration across business networks in context of business processes and the personalization of services to individual preferences.

In order to validate the achievement of these objectives, a testbed infrastructure for SOA4All will be built in the scope of this project. The testbeds created with the SOA4All testbed infrastructure should provide an environment to conduct experiments with and for the two main outcomes of the project – the SOA4All runtime, consisting of a distributed and scalable Service Bus as infrastructural backbone supported by Semantic Spaces for communication and coordination, and the SOA4All Studio, a holistic GUI used to access the different SOA4All components which cover the lifecycle of services, from design-time to run-time and post-mortem analysis. These experiments will focus on functionality tests, and validation of the project results.

In addition, the three use cases developed within the scope of the project will be tested as well, as they will provide a realistic setting for the project.

In addition, the three use cases developed within the scope of the project will be tested as well, as they will provide a realistic setting for the project. One of the Use Cases developed in SOA4All is briefly summarized below, together with an explanation of the kind of services to be used within each scenario:

#### Use Case: End-user Integrated Enterprise Service Delivery Platform

By example of a public administration scenario, this Use Case investigates how existing, heavyweight SOA platforms can interoperate with the open, dynamic, lightweight, and end user driven service platform that is envisioned by SOA4All. The Use Case is working with concrete scenarios implementing the “EC Services Directive” in an informal, narrative way. In these scenarios, several administrative processes need to be completed, triggered by citizens. Following the Services Directive, one public administration takes the responsibility to handle and guide through these processes based on the lightweight process modeling and execution environment developed by SOA4All.

The Use Case uses some SAP Enterprise services, which are not easy to use. The WSDL files are quite complex and some of the data needed is not easy to get, because it is based on SAP code and data structures. Therefore, the current version of services would probably evolve within this case study to simplified WSDL not available now. On the other hand, these services are not services deployed on the web, but in SAP machines. These SAP services are theoretically implemented and deployed on test servers by SAP in SOA4All. So the availability of these services for the testbed is only limited.

It is expected to use third-party services, mainly coming from stub services (simulated services) available from public administration, but these services have not been identified yet.

#### B.11.18 Congestion control - (using Self-Net)

Network operators are focusing on the optimal use and combination of access and core technologies in an attempt to provide consumers network access agnostic services. Heterogeneity is introduced for the benefit of consumers, while network operators undertake the responsibility to manage this complexity within their domain. Congestion Management is a common task that network operators are coming across on a daily basis in such an environment. This Use Case assumes a network infrastructure of WiMAX broadband access provision to consumers. For the Use Case, consumers are using normal (PSTN/ISDN) telephone devices connected to Analogue Telephony Adapters (ATAs) and the WiMAX Subscriber Stations. Consumers initiate VoIP calls, until, after some time, the WiMAX Base Station (BS) identifies congestion on the wireless link and examines potential actions to remedy this. The BS decides to decrease available bandwidth for each flow in order to maintain the current number of voice calls and conveys this decision to ATAs. ATAs then, in order to conform to this decision, switch VoIP calls to a voice codec featuring higher compression. Upon successful completeness, this reconfiguration results in an overall decrease of traffic over the previously congested wireless links with the goal to preserve Quality of Service for the initiated voice calls and maintain the operation of the wireless access network.

#### B.11.19 Context based services - (using C-CAST)

The C-CAST Daily Commuter Use Case focuses on the consumption of content enhanced with contextual information. Specific context information is mixed with content and advertisements, according to the user's location and an environment responding to his individual profile. Commuters are travelling through a town on a train and watch personalised streams, possibly with others who are subscribed to the same stream. Consequently, the commuters are mobile and travel either on the same train (same location) or on a different train (different locations), so they use different access networks and possibly different providers. Vending machines and shops allow for profiled advertising and purchases.