Position Paper: Research Challenges for the Core Platform for the Future Internet

Editors: Michael Boniface, Mike Surridge and Colin Upstill
University of Southampton IT Innovation Centre

1 Introduction

This Position Paper addresses research challenges facing the Future Internet Public Private Partnership (FI PPP) for which Research Centre background and contributions provide key value. It is based on expert contributions from the following European research Centres:

This Paper focuses on challenges for the FI Core Platform which can be met by research results that will be available for implementation in the timeframe of the FI PPP, enabling evaluation by business units by the end of the FI PPP. These include:

- emergent systems engineering and compliance;
- operational risk management;
- turning information into value; and
- socio-economic and user acceptance.

Sections 2 and 3 of this paper review the background and approach to the FI PPP, covering the dual focus on application testbeds and core platform enablers, and clarifying the terminology. Section 4 presents research challenges identified by contributing Research Centres, and explains the requirements for enabling technologies that depend on research results. Section 5 summarises how investment research in generic enablers and capabilities can be leveraged.

2 Background

The Future Internet Public Private Partnership aims to facilitate the development and application of Future Internet technologies, boosting the competitiveness of European Industry, creating new economic opportunities for businesses, and empowering individuals and communities to innovate and benefit from their use.

There is no universally accepted definition of the Future Internet, but stakeholders agree that it will be a socio-technical system comprising Internet-accessible information and services, coupled to the physical environment and human behaviour, and supporting smart applications of societal importance. The Future Internet will therefore become a critical infrastructure for the conduct of business and social interactions, disrupting established business models and value chains (gradually over time if not overnight) but also creating new opportunities. Efforts to create and deploy Future Internet technology must therefore also take account of socio-economic factors including the critical nature of the infrastructure being developed.

The core objective for the FI PPP is therefore to develop generic enabling Future Internet technologies that can provide a platform for the widest possible range of applications. The proposed implementation strategy is to organize activities into vertical applications that build upon horizontal enabling technologies, as shown in Figure 1. The intention is to develop a single Future Internet system with generic enabling technologies that provide the baseline for multiple application-specific test beds.

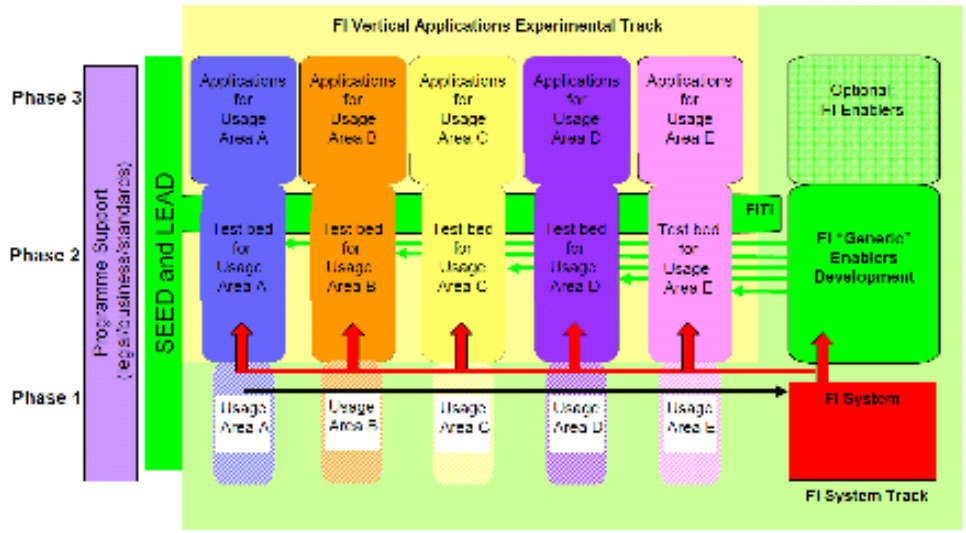

Figure 1. Structure of the FI-PPP

Future Internet applications developed by the PPP are expected to be vertically integrated within sectors, multi-stakeholder (and cross-border), linking information and physical devices, providing facilities for improved understanding, management and operation of socio-economic activities spanning both the digital and physical worlds. A key attribute of Future Internet applications (and the underlying e-infrastructure) is that they should be 'smart', meaning that they are:

- intelligent and able to make choices based on a wide range of information that may be available on the Future Internet;
- flexible and able to dynamically adapt to the needs of (ever-changing) stakeholders by exploiting Future Internet capabilities for rapidly changing connectivity and configuration of facilities and services; and
- efficient and 'green' in their use of physical resources, though improved access to and exploitation of ICT resources.

Future Internet platforms are expected to be largely common across applications and sectors, comprising a collection of enabler technologies providing FI capabilities needed by applications. These platforms should make a significant contribution to the intelligence, flexibility and efficiency of all Future Internet applications, as well as reducing the cost of implementation and operation. Many of the socio-economic aspects (economic viability and sustainability, regulatory and legal compliance, and acceptability to users and the public) should also be addressed in a coherent way via a common infrastructure. Thus FI PPP platforms are also expected to be 'smart', as well as helping to enable 'smartness' in applications built upon them.

Future Internet testbeds will be created by deploying platforms on suitable infrastructures to support applications. These testbeds should leverage previous investments in infrastructure, as well as in platforms capabilities and enablers. Success will depend on the ability to identify and communicate generic capabilities and enablers, to promote broader take-up by longer-term adopters that cross applications and sectors, technology stacks (networks, services, content and things) and vendor-specific implementations and timescales by coupling PPP activities with longer-term research such as that supported by national programmes, the EC framework programme and the EIT KIC 'ICT Labs'. Significant innovation will be necessary to harden current research results, enhance existing commercial software products and fill identified technology gaps to ensure effective evaluation by business units (in contrast to research units) and to provide clear routes to market for products and services.

3 FI PPP Challenges

Based on experience from previous research on disruptive technology (e.g. Grids), several challenges can be anticipated in such a programme. The overall problem is to conduct innovative research to address the socio-economic, technical and application issues, while at the same time providing robust outputs to enable realistic open trials on a large enough scale, and maintaining the potential for commercial exploitation after the research phase.

Technical vs Socio-Economic drivers: History shows that disruptive technology can only have a successful impact through a balanced approach that also addresses socio-economic factors (e.g. the emergence of commercial cloud computing services, c.f. the failure of Grid computing to become ubiquitous). This is true even if huge budgets are used to attempt relatively modest levels of innovation (e.g. the UK's NHS IT Infrastructure project). The PPP will provide a substantial budget for R&D, but this is small compared with the innovation opportunities. It will be tempting to focus resources on technical challenges where research results are most tangible and easily exploited. This must be resisted.

Horizontal vs Vertical drivers: The PPP takes an 'application led' approach to defining and implementing the required technical innovations to create and exploit the Future Internet. This ensures a holistic approach to technical developments within each application sector, and maximises the chance of high-impact exploitation and commercial adoption. However, it will also encourage specialisation in each sector and make it harder to share solutions across the PPP, adding to the overall cost of infrastructure development, and precluding cross-domain business opportunities. A balance must be maintained between near-to-market vertical application drivers and medium-term cross-domain platform utility.

Research vs Development drivers: The PPP is not simply a vehicle for product development. It will have to go well beyond re-badging existing technologies to address research challenges to provide a trustworthy, federated, scalable and converged FI framework capable of addressing smart applications at the technical, socio-economic and regulatory levels. Yet PPP results must also be robust if they are to support large-scale testbeds and lead to short- and medium-term

exploitation opportunities. It will be necessary to balance the need for research and innovation and the need for robust, demonstrable prototypes.

Expoitation vs Efficiency: The PPP aims to provide opportunities for European industry to exploit the Future Internet by providing and/or using innovative technology, but to conduct the necessary research in an efficient way by sharing foreground across the PPP as a whole. However, exploitation may only be possible if the work builds on commercial background that cannot be freely shared, which may inhibit industry partners from implementing complete solutions needed for the Future Internet. The best solution is to ensure that a common baseline technology exists that is standardised or vendor-neutral, allowing individual testbeds and vendors to specialise without creating multiple instances of common capabilities.

Research Focus vs Durability: Focusing on vertically integrated Smart applications will help the PPP to manage some of the above challenges, by limiting the scope of the research and ensuring some level of coherency and interaction between application sectors. However, the PPP should still contribute towards broader (e.g. the Future Internet in rural environments) and longer-term (e.g. virtual living) research challenges. The PPP should produce outputs to support these wider/deeper challenges, as well as starting points for commercial take-up.

4 Core Platform Research Challenges

4.1 Overview

The Internet was originally devised to ensure reliable connectivity in a hostile military setting. However, its use and governance has evolved from that of 30 years ago, and continues to evolve:

- governance objectives have changed from supporting governments and academics to providing an environment in which businesses and citizens may compete as well as cooperate to extract value;
- connectivity (ISPs) is run as a commercial activity rather than by governments and academics, while new service models requiring guaranteed performance have emerged;
- trust between users has reduced dramatically;
- developers are no longer concerned mainly with systemic failure but with the need to develop and maintain emergent applications for potentially competing stakeholders; and
- societal and legal aspects of Internet use (e.g. equality, privacy) are increasingly important for businesses, governments and citizens.

In the Future Internet, these objectives have to be met in the context of technologically converging networks, services, content and devices, and closer coupling of the digital and physical worlds. This will see an increased dependence on distributed information controlled by independent parties, governed by both markets and regulation, so the Internet becomes a critical infrastructure for information exchange, in which the consequences of failure have an impact in the real world.

In this context, we have identified four main areas where FI platform capabilities will be needed, for which significant research challenges have still to be met.

- Emergent systems engineering and compliance: how to design Future Internet systems to meet requirements, given that they will be created and evolve dynamically 'on demand' with no overall designer?
- Operational risk management: how to ensure in real time that systems with no overall controller will operate in a safe and acceptable manner, including interactions with the

physical world, and considering both autonomic and semi-autonomic adaptation processes?

- Turning information into value: how to make information accessible to applications that convert that information into value, and how to preserve this value over long timescales?
- Socio-economic and user acceptance: what platform capabilities are needed to ensure that users and society will accept the Future Internet and use it beneficially?

These capabilities and the associated research challenges are explored in more detail below.

4.2 Emergent systems engineering and compliance

The challenge: to develop enabling technologies for creating systems whose behaviour (or potential misbehaviour) is sufficiently predictable and complies with stakeholder requirements, in an environment where stakeholders have independent goals and may compete (e.g. in a business sense) as well as co-operate to achieve them.

One of the main goals for 'smart' Future Internet applications is to improve the efficiency of physical activities such as energy distribution, transportation and health care delivery by exploiting information shared across organisational and administrative borders to allow non-local optimisation. The resulting critical infrastructures will become increasingly dependent on the FI, so it is important that FI platforms can support a high level of resilience and also correctness.

In such critical infrastructures, tools for ensuring compliance with the relevant regulatory regimes are a non-negotiable checklist item. In the Future Internet, information exchanges will be governed by a combination of business, regulatory and technical measures, including:

- business strategies for responsibility, accountability and governance;
- legal and regulatory mechanisms to safeguard the provision and use of Future Internet services, along with potential government policy initiatives to improve security, trustworthiness and data protection; and
- technical measures to support business and operational approaches to make the Internet more efficient, safe and secure.

Future Internet applications will be extremely flexible compositions of information, devices and services. They will evolve over the long term to meet new requirements, as well as adapting over the short term as new information sources, devices and analysis tools become available. These changes will be driven by multiple stakeholders, responding to their own changing needs as well as changes in the Future Internet environment. Even where Future Internet applications are initially designed by a single authority with the agreement of all stakeholders, over time they will become emergent systems, created by the collective yet independent actions of all stakeholders.

Given these characteristics, a typical Future Internet application will not conform to any a priori system model, and the analysis of its functional and non-functional properties will need to become a dynamic, run-time activity that can be carried out independently by each stakeholder. This is an extremely challenging problem which requires a new approach to system development, spanning the full lifecycle, addressing the socio-economic requirements of stakeholders, and mapping requirements to (dynamically changing) Future Internet platforms and infrastructure. To make the problem manageable, it will be important to develop a 'converged' Future Internet architecture, in which heterogeneous Future Internet components (federated networks, services, content and physical sensors and other devices) are handled in a simplified and uniform way (e.g. by network virtualization). This will involve a convergence between currently disparate architecture

development efforts in the Internet of Things, in the content domain, in the Internet of Services, and in the network domain.

Compliance and security are intersecting but non-identical issues: a secure system may or may not also be compliant, and a compliant system may or may not be secure. Nevertheless, security by design will be an important attribute of FI components. Security-related enablers will play an important role at the platform level in enabling engineers to produce compliant systems. Key enablers are:

- dynamic, cross-domain models of authentication, authorisation and accountability;
- analysis and mitigation of vulnerabilities, including vulnerability to physical and ICT attacks, and also ICT dependency and interdependencies;
- intrusion and (more generally) system change detection and response; and
- auditability of multi-stakeholder, emergent systems.

Other important enablers include well-defined architectures and system modelling approaches capable of handling converged (coupled physical and ICT) systems, standards to ensure interoperability, and a methodology for using FI technologies throughout the life of a system (from design to decommissioning) to achieve system compliance and assurance.

Certification and standardisation of the development and validation and verification methodologies employed should also be considered as a key enabler. The overall approach should extend methodologies such as Control Objectives for Information and Related Technology (COBIT), and design analysis tools, addressing dynamic emergent systems by incorporating dynamic security models. The complexity of the software and technologies will require innovative approaches to the existing testing processes and tools available. Automated tools will be essential in order to conduct robust and repeatable testing and experimental analysis on the large scale projects proposed. An end-to-end robust, adaptive and scalable testing process will have benefit across all the vertical usage areas identified. There may be an opportunity here to develop a new test maturity model – one which is based on existing models (e.g. CMMi, TMMi) but is adaptive and fluid enough to evolve and meet the specific challenges that the innovation of the Future Internet will bring. Standards should build on existing regulations and best-practice for guaranteeing specific systems and devices, such as the Health Information Portability and Accountability Act (HIPAA) or the Payment Card Industry Data Security Standard (PCI-DSS).

4.3 Real-time operational risk management

The challenge: to develop a generic framework for real-time risk management that can link operational monitoring with strategic goals by combing specific risk modelling and analysis techniques with supporting services for automated decision making (policies), risk registers, and risk visualisation.

Critical infrastructures require ICT to support planning and collaborative decision-support processes through access to relevant information, on time and in context. Applications and infrastructure will typically be large-scale, technologically inhomogeneous systems that span organisational and administrative boundaries. The increasing dependence on information governed by independent parties will introduce greater uncertainty into systems (e.g. the delivery of timely and accurate information cannot always be guaranteed, information relevance cannot always be evaluated, etc). The resulting systems will depend on autonomic (e.g. SLA-based) management technology for cost effectiveness. Due to the growing interdependencies between physical and digital worlds, this management will affect real-world interactions as well as the utilisation of ICT resources by applications. Moreover, the presence of uncertainty about

stakeholder behaviour as well as the physical world will make it difficult to ensure that a Future Internet application can meet real-world socio-economic needs. For example, it will be non-trivial to guarantee timely and accurate delivery of energy grid management decisions, given that FI information sources and processing facilities as well as the weather may all be subject to change.

Systems need to adopt an integrated approach to real-time risk management for assessing and dealing with uncertainty considering very large and rapidly changing datasets. The Future Internet infrastructure should provide mechanisms to manage autonomic behaviour in this wider, converged sense, taking account of available Future Internet information. Such mechanisms will provide an underpinning 'smartness' on which Future Internet application developers can build, e.g. opportunistic routing and disruption tolerance. This will require a framework for (real time) assessment of past behaviour and future risks, and adaptation of management policies to manage those risks, based on the self-organisation of many FI system monitoring and management enablers to provide the necessary cognitive capabilities. The types of risks will depend on the application, but may include the behaviour of user communities (e.g. in systems that depend on social networking inputs), physical world events and changes, and changes in the Future Internet infrastructure – networks, services, devices and content.

Techniques will need to be developed to combine pre-processing historical data in batch mode as well as being able to detect, categorise and aggregate incoming events in real time; identify appropriate actions, simulate what might happen if they were taken; and then decide on the best course of action to take and monitor its success. This chain needs to be executed very rapidly and most importantly well within the time scales of change of a community if there is hope of being effective with the actions. These capabilities will exploit autonomic management enablers, but each autonomic system may represent a different stakeholder, and cannot be assumed to cooperate in what is essentially a tussle space. For example, the current business and management models used by interacting ISPs (realised by BGPv4 routing policies and trading limitations) can be expected to break down and lead to instability, especially when other stakeholders such as content owners are also seeking to manage FI system behaviour according to their own interests. The solutions will need to be stakeholder centric, but incorporating new 'risk engine' elements capable of monitoring workflows and balancing risks as well as rewards.

Security risks will play a role, and risk management approaches will be linked to the use of system engineering and compliance (see above). For example, in the Future Internet it should be possible to use agile adaptation strategies to maintain real-time compliance with security and dependability requirements. Also important will be the use of an Extended Dependability Hierarchy approach, in which network and device characteristics are expressed and managed across all levels from the infrastructure to the application. This should be facilitated by an exchange of information between high- and low-level elements, ensuring consistent awareness and treatment of dependability properties across all levels from the infrastructure to the application.

4.4 Turning information into value

The challenge: to develop enabling technologies to discover, access and exploit the wealth of information available from a multitude of (real-time) sources, and to preserve these information assets over very long timescales, meeting societal and regulatory requirements.

The Future Internet will provide unprecedented access to distributed media and other forms of rich content. Making this content searchable and accessible (metadata generation and structuring) will be a major challenge. Users will expect to retrieve what they need, delivered just-in-time. Some users will choose to deal with data, applications, and storage from the cloud as a service infrastructure. Although information search and retrieval will ultimately be an application concern, it is clearly appropriate to support search and retrieval, delivery and (distributed) storage

at the platform level, especially for very large media content objects, for media collections, and for real-time data streams e.g. from sensors. This will have a significant benefit for the cost of developing applications, and for the quality of experience as seen by content consumers.

Metadata generation and standards will be increasingly important, to support content discovery. User-generated metadata will harness the power of social networks to provide alternative views of content and its value that may be more relevant than conventional search and retrieval methods. Standardisation of metadata will ensure interoperability between heterogeneous information sources, not just for consumption of content by applications but also for discovery, composition and long-term preservation.

Ensuring the long term usability of digital information assets is itself a challenge, particularly for highly-regulated societal applications or industrial products with longer lifecycles than the computing systems used for their design, realisation and operation. This is a particular problem that could easily become a barrier for use of the Future Internet, where it will be necessary to preserve federated digital assets across multiple stakeholders. Examples of this include health information systems which may span hospitals and administrations possibly in different EU states, or design information about long-lived products such as aircraft, which may be distributed along the whole supply chain. It will be necessary to preserve the integrity and accessibility of these assets for many years, despite many changes in Future Internet infrastructure and component technologies in that time.

Addressing this challenge requires the integration of efficient, effective and complaint preservation strategies with operational business processes, and supporting these through Future Internet technologies and governance mechanisms. These will need to address the need to protect intellectual property and the privacy of citizens, and allow stakeholders to match their investment to their own needs, while still providing an assurance that the combined set of digital assets will continue to function to the level needed by all stakeholders. The work should build on research such as that in FP7 ICT Challenge 4, but applied to Future Internet technologies and scenarios. This is likely to involve the development of metrics and measurement techniques for long-term federated asset preservation, as well as new strategies and business models for preservation, access and sustainability of collective digital assets, including the use of third party preservation services, data migration/replication strategies, and cost-effective ways to counteract technological obsolescence, etc. Finally, it will be necessary to explore how these approaches can be used to achieve regulatory compliance in specific application sectors.

4.5 Socio-economic and user acceptance

The challenge: to provide platform capabilities to ensure that applications can deliver socio-economic value and can be made open and acceptable to users and society.

The Future Internet must be open to all regardless of their expertise and means of access, and demonstrably fair to participating citizens and businesses. All stakeholders should be able to derive benefit. At the micro level, economic productivity and sustainability of the software industry (e.g services and content) is a major concern. Today, the digital market is focused on the provision of services as a business model from the provider's perspective rather than economic production processes in the traditional sense, i.e. based on revenue. This model of economic exchange is suited to the material economy but does not fit the knowledge economy, which often includes non-monetary exchanges. The assessment of FI ecosystems, underlying value networks, and barriers to productivity will be increasingly important (e.g. technological convergence impacts vendor lock-in, branding formalises lock-in the consciousness of users, and intermediaries constitute barriers to emergent and socially-driven activities). Users will increasingly become part of digital production and distribution with the availability of affordable and accessible FI enabled

tools creating new knowledge-intensive and information rich participative processes that can benefit society as a whole. Governments, citizens and businesses will need to deal with successfully with organisation and cultural changes implied by these new participative models.

Socio-economic and user acceptance will depend to some extent on platform capabilities. These should be addressed in a generic sense at platform level, so that when solutions are found in one domain (e.g. e-health) it is relatively easy to transfer them to other domains (e.g. education and training).

Examples include:

- user and device sensitivity and adaptation, e.g. to match impedance between content streams and the capacity of users to consume them, given their devices, connectivity and expertise;
- human-computer interaction models and modalities, e.g. tools for natural language translation, verbal interfaces for mobile users;
- trading models, mechanisms and standards for provisioning Future Internet systems across borders, e.g. through cloud service marketplaces;
- value models encompassing non-monetary contributions to societal values and well being, and their use in FI management;
- FI health and impact monitoring: supporting measures of overall FI ecosystem efficiency and societal benefits (e.g. levels of accessibility, malicious traffic, trustworthiness,); and
- FI governance models: how government, businesses and citizens together can ensure the FI remains fair, open and socially acceptable – not limited to managing the name and address space.

Unlike the previous capability areas, these are likely to be introduced into platforms following their initial development in application testbeds, possibly in the top-up phase of the PPP. The key contribution of Research Centres will be to provide cross-sectoral analysis to identify transferrable capabilities that can be cost-effectively provided at the platform level instead of being developed in each application.

5 Leveraging Previous Investments

Many generic enablers and capabilities are available from EC and national research programmes (e.g. autonomic networks, cross-layer QoS management, dynamic security and dependability, federated information modelling and distribution, etc). By involving Research Centres to help address the applied research challenges discussed in this Paper, the FI PPP can build on these enablers and capabilities, integrating them into a converged architecture based on open standards, leverage previous investments to create flexible platforms that deliver cost-effective testbeds, and enable socio-economically practical deployment of 'smart' applications on a Future Internet.