

## **Towards a Future Internet of Services**

**A compilation of contributions on Cross Domain Issues  
by members of the post-Bled Software and Services Working Group**

### **Status of this document**

This document is work in progress. It consolidates existing contributions from various members of the post-Bled Software and Services Working Group. The group has regular discussions on the topics addressed in the document in preparation for the Future Internet Assembly (FIA). The document therefore will be updated in the run up to the FIA event organised by the European Commission and hosted by Universidad Politécnica de Madrid on 9 - 10 December 2008 in Madrid.

The document covers the following four cross domain issues:

- Architectures and Infrastructures
- Management and Governance
- Trust at Scale and High Granularity
- Lifecycle engineering for Future Internet Applications.

## Architecture and Infrastructure as means of convergence for the Future Internet

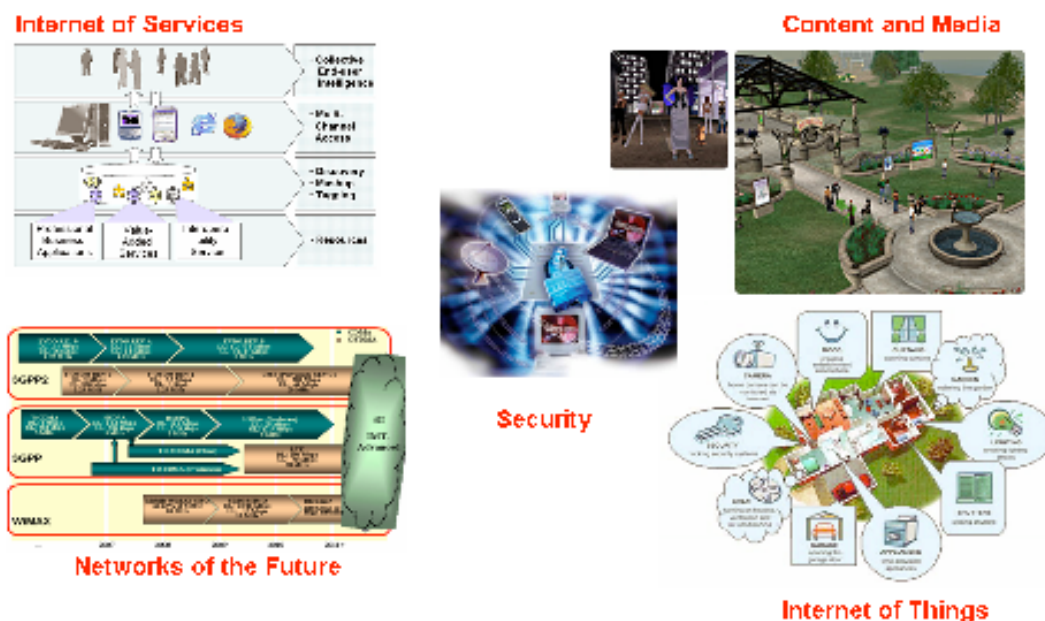
**Contributor: Stefano De Panfilis, Engineering**

The Internet has changed society and is increasingly shaping social communities and business interactions. The Internet has evolved from the largely static Information Super Highway of the 90's to a critical infrastructure supporting all aspects of life in the 21st Century. Internet applications increasingly require a combination of capabilities from traditionally separate technology domains to deliver the flexibility and dependability required by users.

The next decade will introduce remarkable progress in a wide range of pervasive technologies culminating in the introduction of the Future Internet. Beyond technological aspects, the Future Internet is likely to have a profound effect on our society, from a societal, organisational or business perspective.

The EIFFEL Think-Tank report (<http://www.future-internet.eu>), which was released in December 2006, defines the Future Internet to be a global, open platform and identifies a number of drivers for the next generation Internet, such as emphasis on mobility, the anticipated changes in scale of connected devices, increases in bandwidth, increase in digitised media, increasing importance of security and evolution of services to more adaptability, and awareness of user context and preferences.

To achieve the objectives of the Future Internet at least five communities need to converge and cooperate: Networks, Internet of Things, Internet of Services, Content and Media, and Security. The objective is to identify, understand and resolve convergence challenges between these technology domains and to give Europe the lead edge over other competing global initiatives.



**Figure 1 Future Internet: converging domains.**

The purpose of this Working Group is to identify architectural and infrastructural concerns that will make the Future Internet a reality. Architecture and Infrastructure are the foundations on which Future Internet applications will be developed.

Typical characteristics of the architecture and infrastructure of the Future Internet in terms of the five communities identified in Figure-1 are:

- Networks *<to be filled by the networking community>*
- Internet of Things *<to be filled by the relevant community>*
- Internet of Services:
  - Service architecture and infrastructures aim at standardise the way services are built, offered and consumed. There is a clear need for introducing a Reference Service Architecture for materializing the Future Internet.
  - The Reference Architecture should be defined in order to allow any business domain, business of any size, and technology to engage in a collaborative and dynamic society.
  - Implementing a Reference Architecture is not a single shot effort but an incremental process with should involve many participants both from academia and industry.
- Content and Media *<to be filled by the relevant community>*
- Security *<to be filled by the relevant community>*

With the full involvement of the FP7 Projects, and particularly of those who are members of the Future Internet Assembly, this Working Group will proceed step by step starting from identifying relevant results of each project, and by then clustering the results in common and shared items which will be the initial components of the Architecture and Infrastructure of Future Internet.

## **Management and Governance**

**Contributor: Alex Galis, UCL**

**Problems:** Many identified bottlenecks of the current Internet of Networks and Services including ossification and non-imbedded manageability.

### **Systemic Research Scope:**

- Define cross-domain (i.e. networks, services, content) and cross-layers cooperative FI system design for integrated management functionality, including: system lifecycle- , monitoring- , (re)configuration- , optimisation- , organisation- , performance- , adaptation- , context- , semantic- , security- , composition- , assurance- , negotiation-, repository- , SLA- , QoS- , billing-management and self-management
- Imbedding management functionality in all FI systems (i.e. InNetworks management, InServices management, InContent management)
- Dynamic deployment of (new)management functionality without interruption of FI systems' and services' operation (i.e. Plug-and-Play, UnPlug-and-Play, programmability)
- Orchestration and integration of management functionality
- Define Self-contextualisation for FI systems, services and resources
- Define FI Autonomicity and Self-awareness
- Define integrated and flexible usage of heterogeneous and assumable resources for energy, networking, computation, storage, content, mobility, etc
- Minimise life-cycle FI system costs, minimise energy footprint

**Approach:** Leverage and influence effort in existing FP7 projects

### **Expected Results:**

- Common position(s) on the necessary and prioritised research challenges and reference configurations; identify key directions for future research projects
- Reference management architectures and infrastructures; alternate paths
- Define management of virtualised FI systems and services

## **Trust at Scale in the Future Internet**

**Contributors: Nick Wainwright and Neil Dunbar, HP Laboratories, Europe**

In the Future Internet of pervasive services and devices people will interact every day of our daily lives with hundreds of independent computing devices and services. It is vital that and that people and organisations develop the ‘trust’ in these devices, systems and services that will give them the confidence to participate in the future internet, and make use of the many valuable services available to them.

Yet in a world of billions of connected devices, the traditional models used to provide trust start to break down. Hierarchical models of identity and authority lose cohesion, and the axioms upon which one would build a calculus of security assertions become too vague in order to construct a sufficiently useful policy enforcement method. Similarly, peer-to-peer ad-hoc trust communities cannot guarantee that the assertions made within their community are compatible with others, and thus transitivity of assertions is lost.

Thus, we believe that new models of identity acquisition and behaviour control need to be developed in order to make the future internet a trustworthy space. Such models must address vastly divergent computing, storage and communications capabilities, as well as far more nuanced concepts of multi-party ownership, transient control and partial failure. The models must encompass not merely the traditional guarantees of confidentiality, integrity, availability and privacy, but must embed emergent qualities such as data usage scope, transitive rights revocation and fair use of digital services. Failure to do so will limit the investment that European citizens are willing to make in a pervasively connected Future Internet (FI). We propose to research and construct new types of device platform, whose agents compute from a data governance perspective, rather than a hardware utilization one, actively attesting to their behaviour at all times and being able to account for decisions made on behalf of the data owner or his/her other agents.

In addition to the problems posed to end users by the scale of the trust problems, a parallel set of management problems emerge to plague system administrators, network operators and policy enforcers. Access control and auditing requirements are often tightly tied to ownership domains. Such persistent control of domains of well known assets cannot be relied up on in a large scale pervasive future internet. Indeed, it is arguable that such a level of control is not even present now, and that many system administrators choose to turn a blind eye to participating devices not officially within their asset base.

We believe that the administration of services in the future internet world will need more far more flexible and powerful models for service provision, data governance and device configuration if the costs of building and running a high quality internet space are to be affordable. We propose to construct security management frameworks based upon current research in Mobile Ad-Hoc Networks (MANET), planetary scale computing efforts, trusted virtualization and in order to create appliance networks which largely automate the deployment and auditing life cycles of both hardware and software according to well specified policy requirements, which operate across the differing domains of personal, business and governmental computing sectors.

## Lifecycle engineering for Future Internet Applications

**Contributor: Mike Boniface, IT Innovation Centre**

Many challenges exist when we look at the interplay between content engineering, service engineering and network engineering lifecycles. For example, content and service lifecycles are changing radically and we need to understand the impact. The challenge is actually how these things can be kept separate, i.e. move away from traditional models where they are all locked together (for example, consider television and how the way that content is commissioned and shot is actually a result of the way people will view it on a TV set and the way it will get to them over the airwaves - the device, channel and content are all connected). How can content be engineered so it remains usable when the devices and networks used to produce, consume and distribute it are transient? Furthermore, how can content be engineered when these devices and channels may not even exist yet? This is about levels of abstraction and having flexible processes and services - just in the same way that we use these ideas in software and systems engineering - and using these techniques to avoid solutions that are brittle and don't meet user needs. The key here is to begin dialog on vision and operation models for the future internet and a cross-domain dialog from engineering lifecycle perspective provides a good foundation for that discussion. Therefore, the proposed topic is to look at a broader notion of engineering applied to services, content, things etc. that addresses the interplay between them through the various lifecycles they have and seeks to decouple these lifecycles from each other.

**Lifecycle Engineering for Future Internet Applications**

- Critical for development and operation of Future Internet applications
- Helps to clarify vision, future operation models and applications
- Complex and changing rapidly in all technical domains
- Interplay between application, system, and stakeholder lifecycles needs analysis

**Example: Content and Service Lifecycles**

**A few content trends**

- Creation of content without knowing the audience
- 70% of content is UGC today
- More content generated in 2 years than the entirety of human history
- Examples of data always online

**Impact on service Infrastructure**

- How do service SLAs need to change (metrics, guarantees, etc) to support content lifecycles (e.g. preservation) ?
- How can we keep content safe for 20 years when services and technology are transient?
- How do we ensure that service and content adaptability approaches are cohesive?