Position Paper: *Management and Service-aware Networking Architectures (MANA) for Future Internets*

Draft 5.0 - 3rd December 2008

Editors: Alex Galis (UCL), Marcus Brunner (NEC), Henrik Abramowicz (ERICSSON); Email: <fia-mana@ee.ucl.ac.uk>

1. Introduction

Future Internets – Overview and Context

The "Future Internet" research and development threads have recently been gaining momentum all over the world. On 31 March 2008, the Future Internet Assembly was kicked off at the Bled conference organised by the EC and the Slovenian EU Presidency as the means to enable fundamental and systemic innovation in networking and services for the realization of the Future Internet.

The current Internet has been founded on a basic architectural premise, that is: a simple network service can be used as a universal means to interconnect intelligent end systems. The current Internet is centred on the network layer being capable of dynamically selecting a path from the originating source of a packet to its ultimate destination, with no guarantees of packet delivery or traffic characteristics. The end-to-end argument has served to maintain the desire for this simplicity. The continuation of simplicity in the network has pushed complexity into the endpoints, thus allowing the Internet to reach an impressive scale in terms of inter-connected devices. However, while the scale has not yet reached its limits, the growth of functionality and the growth of size have both slowed down It is now a common belief that the current Internet is reaching both its architectural capability and its capacity limits (i.e. addressing, reachability, new demands on QoS, Service /Application provisioning, etc).

Although the current Internet, as a ubiquitous and universal means for communication and computation, has been extraordinarily successful, there are many challenges some with fundamental aspects. Many of these aspects could not have been foreseen when the first parts of the Internet where built, but these need to be addressed now. The very success of the Internet is now creating obstacles to future innovation of both the networking technology that lies at the Internet's core and the services that use it. In addition, the ossification of the Internet makes the introduction and deployment of new network technologies and services very difficult and very costly.

We are faced with an Internet that is good at delivering packets, but shows a level of inflexibility at network layer and a lack of built-in facilities to support any non-basic functionality. The aspects, which we consider to be missing, are:

- Mobility of networks, services, and devices
- Facilities to support Quality of Service and SLAs
- Trust Management and Security; Privacy and data-protection mechanisms of distributed data
- An adequate addressing scheme, where identity and location are not embedded in the same address
- Inherent network management functionality, specifically self-management functionality
- Facilities for the large scale provisioning and deployment of both services and management services; support for higher integration between services and networks
- Facilities for the addition of new functionality, including capability for activating a new service on-demand, network functionality, or protocol (i.e. addressing the ossification bottleneck)
- Facilities for orchestration of security, reliability, robustness, mobility, context, service support, and management for both the communication resources and the services' resources
- Business neutral way of networking through separation of mechanism and policy
- Energy awareness

The current trend for networks is that they are becoming service-aware. Service awareness itself has many aspects, including the delivery of content and service logic, fulfilment of business and other service characteristics such as Quality of Service (QoS) and Service Level Agreements (SLA) and the optimisation of the network resources during the service delivery. The design of Networks and Services is moving forward to include higher levels of automation, autonomicity, including self-management. Conversely, services themselves are becoming network-aware. Networking-awareness means that services are executed and managed within network execution environments and that both services and network resources can be managed uniformly in an integrated way.

Scope

This position paper identifies the research orientation together with the key challenges for the capabilities and the systems in the *Management and Service-aware Networking Architectures (MANA)* part of the Future Internet.

In order to achieve the objective of being service-aware and network-aware (i.e. servicing and networking resources must be aware of relevant environment conditions as well of there own state: self-aware), and to overcome the ossification of the current Internet, this position paper envisages various novel solutions for the Future Internet. The Future Internets must be built as service-aware and self-aware federated networks, which provide built-in and orchestrated aspects such as: context, reliability, robustness, mobility, security, service support, and self-management of the communication resources and the services. Such aspects suggest a transition from a service-agnostic Internet to a new service-aware and self-aware Internet, in which self-awareness is serving the purpose of communication and computation by means of enhanced in-network and in-service decisions enabled by aware-sensing.

In this paper we also envisage capabilities spanning a range of technologies, including:

- Mobile, wireless and high function network core, edges and service nodes.
- Service-aware scalable and robust networking architectures, including:
    - Connectivity-to-network, network-to-network services, network service-to-service computing clouds and other service-oriented infrastructures.
    - Cross-domain interoperability and deployment.
    - Optimal orchestration of available resources
- Management systems covering FCAPS functionality, including self-awareness and self-management (i.e. all self-X functions).

2. MANA Research Orientation / Grand Challenges:

This section summarises the important research challenges and opportunities for the definition and design of *Management and Service-aware Networking Architectures* for Future Internets. We present the main areas where we feel research into the definitions and the design is required, as a starting point.

I. General Capabilities:

- Mobility anywhere and anytime.
- Connectivity anywhere and anytime.
- Adaptability anywhere and anytime.
- Resiliency and survivability anywhere and anytime.
- Robustness and stability anywhere and anytime.
- Accountability anywhere and anytime.
- Evolvability anywhere and anytime.
- Scalability.
- Reliability.
- Trust and security
- Multi-domains.
- Green communication and servicing.

II. Infrastructures Capabilities:

- Computing, networking, and storage elements represent the components of the MANA infrastructure. The MANA infrastructure consists of evolving and expandable clusters of computing, networking, and storage elements (e.g. deployed both on network systems, nodes and on Users' devices), which support the configuration and deployment of any resources (i.e. including virtual resources) of both networks and services. Infrastructure components include:
  - Core Nodes for the provisioning of high-speed, high volume traffic flows for data processing functions, including flexible control and management capabilities.
  - Edge Nodes and Service Nodes for the programmatic provisioning of the transport, computational, storage and content resources needed to deploy wide-area services and new network functionality, including programmability of the network forwarding functions and flexible control and management capabilities.
  - Mobile Nodes and Wireless Nodes for the programmatic provisioning of the communication and computational resources needed to deploy wide-area services and new network functionality within a wireless or mobile network, including programmability of the network functions and flexible control and management capabilities.
  - Programmable data and service centres for the provisioning of networking computational resources for service clouds.
  - All the above in the radically distributed and composable manner
- Ubiquitous Connectivity, Computation, Storage and Content infrastructures, together with the architectures, resources, self-management, and controls of such resources, including the assessment of infrastructure adaptations based on context-awareness.
- New globally accessible Infrastructure Services, including Information-centric and Context-centric networks.

III. Control and Elasticity Capabilities:

- Uniform control frameworks for the Future Internet. These have to be scalable and dynamic, yet be serving diverse operational and business requirements.
- Uniform mobility frameworks for Future Internet.
- New naming frameworks, including Identity / Location splits and support for addressing information or context objects and services.
- New tuneable protocols for different layers of the protocol stack in support of cleaner cross-layer interaction and dynamic service composition.
- Flexible and cost effective operations of service platforms over core and edge transport networks.
- Mechanisms and interfaces to accommodate the conflicting interests of stakeholders in the Future Internet architecture.
- Interworking with the existing Internet.

IV. Accountability Capabilities:

Whilst the need for accountability was known in the every early days of the Internet, it was safely omitted from the initial deployment stages because each player knew the others, and all understood the limitation of the technical platform they were creating. Today the network is built from thousands of smaller networks, and they are supporting a much wider range of uses. This has led to tension and tussle between all the different players. We aim at an "Accountable" Internet - where users are held accountable for any misbehaviour or congestion they cause - hence they are held accountable for their impact on others. As such we need an open delivery infrastructure that can accommodate innovation both at the network and service layer, including the aim to integrate both the technical and socio-economic aspects into a single solution.

- Cross layer optimization – network, transport and service layers - to enhance session-less application driven QoS approaches

- Resource Pooling - Cost effective way for the Internet to achieve high network utilization and secure future innovation where separate network resources behave like a single large pooled resource.
- Multi Transport Congestion Protocol - that combines multipath routing with congestion control and allows traffic to move away from congested links.
- Enhanced Service Control - Enable increased control to the application when applications are best placed to choose the best path for transmission (e.g. low cost path) and manage mobility and multi-homing.
- Enhance Information exposure - Traffic carries info about its resource usage in such a way that the network can monitor the cost (e.g. congestion) of carrying a specific packet but also the application can select the most convenient path to send specific traffic.
- Lightweight Control Architecture - Avoid locating any mechanisms at network resources themselves for resolving usage conflicts with most of policing and management located at the 'enforcement point' – network ingress where customer attaches.
- Separate policy and mechanisms - Need common mechanisms across the infrastructure to control resource usage while the policy can be left under the control of the various stakeholders.

V. Virtualisation of Resources and Service Computing Clouds Capabilities:

- Ubiquitous Virtual Resources with integrated self-management of those resources. This allows for the integrated and flexible usage of heterogeneous and assumable virtual resources for networking, for computation, for storage, for content, and for mobility.
- Virtual assurable resources, which do not necessarily correspond to administrative, topological or geographical domains and which would take into account concerns such as confidentiality, availability, integrity, and safety; are used to enable collaborative groups of consumers to exchange information in pursuit of shared interests, services or business processes.
- Security concerns related to the use of virtual resource.
- Virtual resource-facing services enabling flexible usage of the physical resources.
- Real-time service computing clouds and virtual-private service clouds, integrating the necessary storage, networking and service resources.
- Ubiquitous light-weight virtual channels for integrating an Internet of Things into a service-aware network infrastructure.
- Overlays for enabling decentralized component interactions and for the provisioning of virtualisation of the infrastructure resources; overlays for creating a topology of nodes for the interactions of different components.

VI. Self-management Capabilities:

- Cross-domain management functions, for networks, services, content, together with the design of cooperative systems providing integrated management functionality of system lifecycle, self-functionality, SLA, and QoS.
- Embedded management functionality in all Future Internet systems, such as in-infrastructure management, in-network management, in-service management, and in-content management.
- Mechanisms for dynamic deployment of new management functionality without interruption of running Future Internet systems. The operations required: Plug-and-Play, Unplug-and-Play, and (re)programmability of the forwarding and control planes.
- Mechanisms for dynamic deployment of measuring and monitoring probes for services' and network' behaviours, including traffic. SLA-aware sensing and continuous monitoring of systems' adaptations. Use of monitoring services in support of the self-management functionality.
- Mechanisms for conflict and integrity-issues detection and resolution across multiple self-management functions.
- Mechanisms, tools and methodology for the verification and assurance of different self-capabilities that are guiding systems and their adaptations correctly.
- Mechanisms for allocation and negotiation of different resources.

- Increased level of self-awareness, self-knowledge, self-assessment and self-management capabilities for all Future Internet systems, services, and resources.
- Increased level of self-adaptation and self-composition of resources to achieve effective, autonomic and controllable behaviour.
- Increased level of self-contextualisation and context-awareness for network and service systems and resources.
- Increased level of resource management, including discovery, configuration, deployment, utilization, control and maintenance.
- Self-awareness capabilities to support system-level objectives of minimizing system life-cycle costs and energy footprints.
- Orchestration and integration of management functions. Service driven dynamic orchestration.
- Capabilities for the control relationships between Self-Management and Self -Governance of the Future Internet.

The overall capabilities of the Future Internet architecture are depicted in Figure 1.

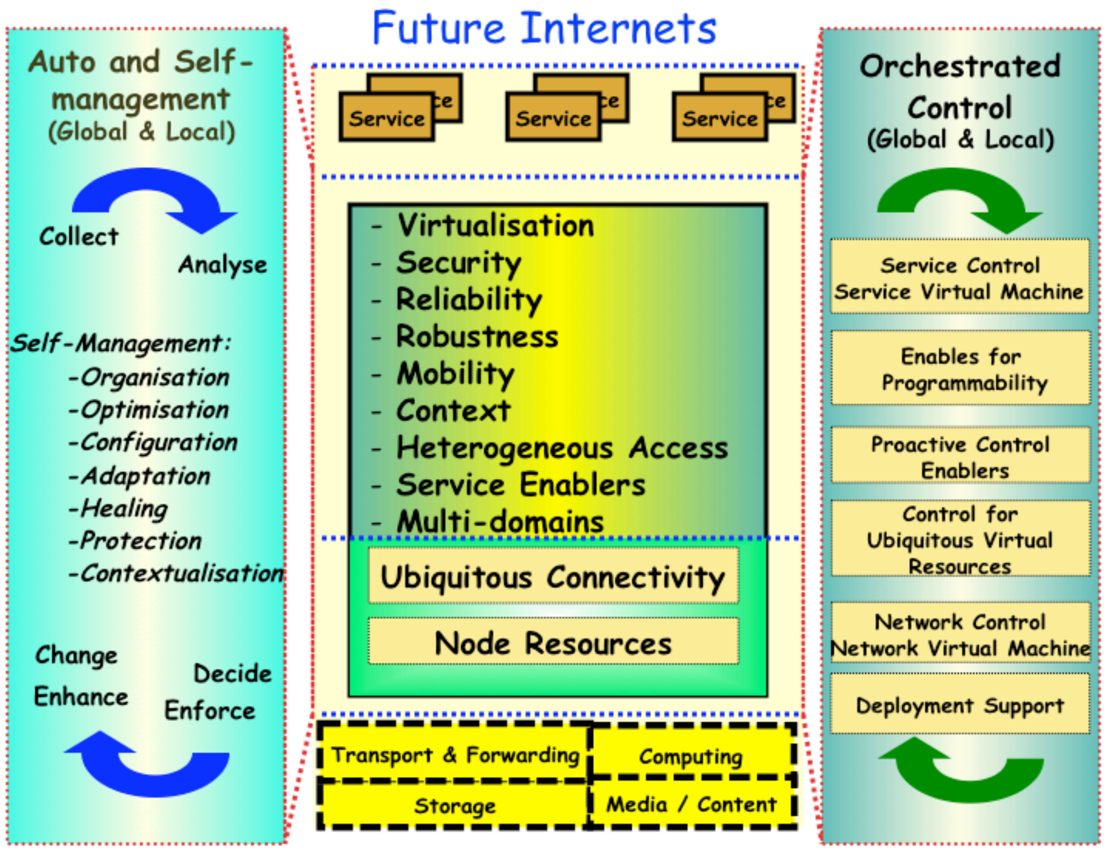

Figure 1 – Future Internet Capabilities

4. Possible integration paths towards the Future Internet

The following steps for progress are envisaged:

- Analysis of the problems and bottlenecks of the current Internet, leading to a basis of research papers.
- Proposals for MANA architectures and systems for evolutionary and clean slate approaches aligned with visions of other cross-domain topics or FP7 projects.
- Proposals for Evaluation, Demonstrations and Interoperability.

5. Roadmap and Milestones

- Preparation of the "Management and Service-aware Networking Architectures" Position Paper - version for discussion at the FIA meeting 9th December 08; final version 15th December 2008.
- Proposition of use cases for use of experimental facilities – initial version for discussion at the FIA meeting in Madrid 9th Dec08; final version 15th Dec 2008.
- System Functions and Requirements (SFR) Paper on "Management and Service-aware Networking Architectures for Future Internet": TOC/Scope prior to FIA Madrid, identify a set of authors at the FIA meeting in Madrid, first draft January 2009 --> to be presented at the next FIA in Prague, 11-13 May 2009.
- Workshop during 2009 to progress the System Functions and Requirements (SFR) Paper on "Management and Service-aware Networking Architectures for Future Internet.

6. References

State of the Play / State of the Art / White papers

7. Contributors

Name	Company & E-mail
Henrik Abramowicz	Ericsson Research, Sweden; <henrik.abramowicz@ericsson.com>
Nancy Alonistioti	University of Athens, Greece; <nancy@di.uoa.gr>
Andreas Aurelius	Acreo AB, Sweden; <Andreas.Aurelius@acreo.se>
Marcus Brunner	NEC Research, Germany; <Brunner@nw.neclab.eu>
Moiso Corrado	Telecom Italia, Future Center, Italy; <corrado.moiso@telecomitalia.it>
Paolo Bellavista	Università degli Studi di Bologna, Italy; <paolo.bellavista@unibo.it>
Stuart Clayman	University College London, United Kingdom; <s.clayman@ee.ucl.ac.uk>
Ranganai Chaparadza	Faunhofer FOKUS, Germany; <Ranganai.Chaparadza@fokus.fraunhofer.de>
Piero Castoldi	Scuola Superiore Sant'Anna, Italy; <castoldi@sssup.it>
Franco Callegati	Scuola Superiore Sant'Anna, Italy;<franco.callegati@unibo.it>
Alan Davy	Waterford Institute of Technology, Ireland; <adavy@tssg.org>
Steven Davy	Waterford Institute of Technology, Ireland; <sdavy@tssg.org>
Henk Eertink	Telematica Institut, The Netherland; <Henk.Eertink@telin.nl>
Alex Galis	University College London, United Kingdom; <a.galis@ee.ucl.ac.uk>
Jose Jimenez	Telefonica, Spain; <jimenez@tid.es>
Michael Kleis	Fraunhofer FOKUS, Germany; <Michael.Kleis@fokus.fraunhofer.de>
Theo Kanter	Mid-Sweden University, Sweden; <theo.kanter@miun.se>
Apostolis Kousaridas	University of Athens, Greece; <akousar@di.uoa.gr>
Miguel Ponce de Leon	Waterford Institute of Technology, Ireland; <miguelpdl@tssg.org>
Antonio Manzalini	Telecom Italia, Italy; <antonio.manzalini@telecomitalia.it>
Panagis Magdalinos	University of Athens, Greece; <panagis@di.uoa.gr>
Hermann de Meer	University of Passau, Germany; <hermann.demeer@uni-passau.de>
Angel Ferreiro Olivo	Telefonica, Spain; <olivo@tid.es>
Costas Polychronopoulos	University of Athens, Greece; <cpoly@di.uoa.gr>
Eleni Patouni	University of Athens, Greece; <elenip@di.uoa.gr>,
John Strassner	Waterford Institute of Technology, Ireland; <jstrassner@tssg.org>
Antonio F. Gómez Skarmeta	Universidad de Murcia, Spain; <skarmeta@um.es>
Mikhail Smirnov	Faunhofer FOKUS, Germany; <mikhail.smirnov@fokus.fraunhofer.de>
Rao Sathya	Telescom, Switzerland; <rao@telscom.ch>
Patrick Strating	Telematica Institut, The Netherland; <Patrick.Strating@telin.nl>
Andrea Soppera	British Telecom, United Kingdom; <andrea.2.soppera@bt.com>
Makis Stamatelatos	University of Athens, Greece; <makiss@di.uoa.gr>
Tanja Zseby	Faunhofer FOKUS, Germany; < Tanja.Zseby@fokus.fraunhofer.de>