Position Paper of FP7 Future Internet Cluster

November 2014

Contributors:

Aiko Pras, Anastasios Kourtis, Andreas Petlund, Andy Edmonds, Antonio Cimmino, Burkhard Stiller, Carmen Guerrero, Corinna Schmitt, David Griffin, David Ros, Dimitri Staessens, Eleni Trouva, Elio Salvadori, Elisa Rojas, Filip De Turck, George Xilouris, Gianmarco Panza, Holger Karl, Lajos Hanzo, Lorenzo Iacobelli, Marco Mellia, Mayutan Arumaithurai, Miguel Ponce de Leon, Miguel Rio, Olivier Festor, Özgü Alay, Peter Willis, Philip Eardley, Pieter Simoens, Stein Gjessing, Thomas Michael Bohnert, Wouter Tavernier, Yan Zhang, and Yannick Le Louédec.

Editors:

Filip De Turck, Thomas Michael Bohnert, Antonio Cimmino

Involved FP7 projects of Future Internet Cluster:

BUTLER, CONCERTO, eCOUSIN, FLAMINGO, FUSION, GreenICN, LEONE, MCN, mPlane, NetIDE, ONE, PACE, PRISTINE, RITE, ROMEO, SmartenIT, T-NOVA, Trilogy 2, UNIFY

1. Introduction

This position paper is written by the FP7 Future Internet Cluster based on a request of the European Commission for suggestions to shape the work program in 2016-2017 of the H2020 calls on Smart Networks and Novel Architectures. The position paper outlines the topics, which the FP7 Future Internet Cluster strongly believes will be of key importance for the future. Focus is on topics where in-depth research is needed during the next years, which are currently not covered in ongoing projects, and which will allow European industry and academia to excel during the next years. Each Future Internet Cluster project was kindly invited to contribute to this position paper, the methodology for establishing the position paper is detailed below in Section 2.

The following 19 ongoing FP7 Future Internet projects participated (in alphabetical order): BUTLER [1], CONCERTO [2], eCOUSIN [3], FLAMINGO [4], FUSION [5], GreenICN [6], LEONE [7], MCN [8], mPlane [9], NetIDE [10], ONE [11], PACE [12], PRISTINE [13], RITE [14], ROMEO [15], SmartenIT [16], T-NOVA [17], Trilogy 2 [18], UNIFY [19].

Each of the contributors listed above belongs to at least one or multiple of these mentioned ongoing FP7 Future Internet cluster projects. All contributors clearly mentioned their current achievements and focus, their planned contributions by the end of their project, and the topics to be addressed in the future. All contributors were particularly encouraged to provide topic descriptions in the domain of the project they are currently involved in (as they are considered to be the experts in this domain and have a clear view on which achievements the current ongoing projects will bring). However, their input to other domains was welcomed as well.

2. Methodology

In order to realize this Future Internet position paper, the following actions were taken (in each of the steps listed below the project coordinators of the 19 ongoing Future Internet Cluster projects were involved):

1. All project coordinators of ongoing Future Internet Cluster projects were invited on September 1st 2014 to contribute: a clear template was sent to them, requesting descriptions of future topics in 10-15 lines and also requesting to mention the main achievements of their ongoing project if the topics build further on this ongoing project, and also to put their ideas for future directions into the context of the achievements of the past/current project.
2. Submission deadline for the above input was September 22nd, 2014 and a few reminders were sent.
3. The received inputs were presented in a dedicated time slot on September 30th, 2014 during the EU Stakeholders Consultation Workshop in Brussels. Over 200 experts participated to the stakeholder consultation work, a very interesting event to learn viewpoints, visions and network with European researchers and project coordinators.
4. The feedback from the participants of the EU Stakeholders Consultation Workshop and additional input from the FP7 Future Internet Cluster project coordinators was taken into account determining the topics and descriptions as provided below.

5. All Future Internet Cluster projects were invited for the Net Tech Future Coordination meeting (formerly called 'concertation meeting') on October 23th, 2014 in Brussels. The meeting was composed of a Future Internet Cluster meeting before noon and a plenary session with all clusters in the afternoon.
6. The agenda of the Future Internet Cluster meeting was dedicated entirely to presenting the current and planned achievements of the ongoing projects and determining and discussing the topics to be suggested for future calls in the area of Smart Networks and Novel Architectures. The meeting was well attended and interesting discussions took place, leading to the list of topics as presented below.
7. During the plenary session of the Net Tech Future Coordination meeting on October 23th, 2014, the identified topics below were presented.
8. A completed version of the position paper was sent to the project coordinators and participants of the above meetings on November 12, 2014 and 1 week was given for final feedback.
9. The final received inputs were taken into account and the position paper was finally delivered to the EC officers on November 21, 2014 and all contributors received the submitted version.

3. Identified Important Topics

Based on the methodology described above, the following 7 important topics have been identified. Possible prioritization of those topics was deliberately not discussed. The Future Internet Cluster suggests these topics to be addressed in future Horizon 2020 calls on Smart Networks and Novel Architectures:

1. Advanced content delivery systems
2. Measurement-based management
3. SDN-based systems and applications
4. Advanced NFV-based systems
5. QoE-centric management
6. Advanced security for smart networks
7. Advanced Internet architectures

4. Detailed Description per Topic

The detailed description of each of those seven topics refines them under the assumption that functionality as well as mechanisms will be beneficial for future systems. As such potential benefits and advantages of the deployment of those are summarized.

4.1 Advanced content delivery systems

Current projects in the Future Internet Cluster focus on crawling and data analysis of online social networks with the aim to optimize content delivery, reduce network costs, improve quality of experience, improve marketing and advertisements, and protect users' privacy. Due to the expected 10-100x increase in the number of devices and a 100-5000x increase in traffic over the next 10 years, primarily driven by Internet of Things and media applications, the following important topics to be addressed in future projects are put forward: increased scalability, flexibility as well as cost and energy efficiency. For these reasons, the Future Internet Cluster believes priorities for the next work program on advanced content delivery systems should include the following topics.

First, research for advanced Internet architectures tightly integrating network functionalities and content-related services should be performed, addressing the coordination among layers and adaptation according to the network conditions, especially in Internet-of-Things scenarios where communication bandwidth can be limited in some parts of the network while the aggregated traffic will be huge in the core network and of fundamentally different nature than what we observe today and on what our models are currently built. Observability of network performance at large for improving the performance of inter-domain traffic and for optimizing service orchestration and content delivery can be an important property of such a novel content delivery system.

Second, fixed-mobile convergence scenarios for content delivery should be considered in much more detail, encompassing edge computing and content distribution. An important topic is the support for massively distributed service deployment at the network edge for highly interactive, extremely low-latency, real-time services (e.g. virtual/augmented reality and safety-critical systems such as self-driving vehicles and medical applications), requiring a joint management of service orchestration and the underlying network.

Third, advanced large-scale content dissemination systems exploiting social information and dynamic adaptation based on user preferences are very important. The option of using available resources in already existing network nodes in the delivery path can be explored (together with advanced caching mechanisms) in order to facilitate intelligent multimedia data delivery while relieving the Content Delivery Network (CDN) load and reducing energy consumption of CDN data centers.

4.2 Measurement-based management

Two current projects address the design of an efficient measurement plane for the Future Internet. Massive amounts of measurement data are gathered. Topics for future projects include data analytics techniques to determine root cause of failures, measurement of extra contextual parameters influencing the users' perception and automated generation of actions based on the massive amounts of gathered data.

A second interesting area is to not only measure QoS (Quality-of-Service) or QoE (Quality-of-Experience) metrics, but also measure whether policies (such as privacy of users when using web-based and cloud-based applications, regulation of cloud systems, price discrimination, verification of information) are

respected and to which extent. This would allow for third-party verification of the services and implement verification and trust in the Future Internet.

Third, network performance metric observability at large (expose and use monitoring data of network, service, and application level parameters, including on-line social networks) would be beneficial for improving the performance of inter-domain traffic and for optimizing service orchestration and content delivery.

Finally, measurement-based management in large scale Internet-of-Things environments and mobile broadband networks is considered to be very important for future projects. There is a strong need for objective data about stability and performance of mobile broadband networks, and for tools to rigorously and scientifically assess their performance. In particular, it is important to measure and understand the quality as experienced by the end user. Such information is very valuable for many parties including operators, regulators and policy makers, consumers and society at large, businesses whose services depend on mobile broadband networks, researchers and innovators. Therefore, large-scale measurements in these environments are going to be of utmost importance, as a tool both for network performance management and for the design of a robust future Internet. This will require the design, development and deployment of new large-scale techniques, tools and open measurement platforms and testbeds for large scale Internet-of-Things environments and mobile broadband networks, with a focus on end-user QoE. The serving cloud computing infrastructure should also be taken into account during the design phase of tools and monitoring engines.

4.3 SDN-based systems and applications

Recently, leading equipment providers in the network infrastructure market launched the first software-enabled appliances that support network virtualization capabilities. The main advantage of Software Defined Networking (SDN) is the separation of network control and forwarding and the fact that it allows flexible management of the network resources. OpenFlow is currently the most prominent SDN-based communication protocol. Future research should be devoted to design SDN-based systems and applications, which exploit the potential of SDN much more than the current OpenFlow-based systems. Research in the following areas is proposed by the Future Internet cluster.

First, the architects and developers of SDN-based network infrastructure need agile and productive design environments (i.e. SDKs - Software Development Kits) comparable to existing ICT software technology development. These library-based SDKs will act as mediator to existing and future SDN control planes and will be used by orchestrator services, frameworks and tools for Virtualized Network Functions (VNFs). Future research should propose new or existing suitable languages, library structures, modularity and universal design environments suitable for both the phase of design and run-time (cloud orchestration). This will enable researchers and engineers experiment with new ways of creating new or optimized existing protocol stacks. The research should further extend SDN from the cloud infrastructure to the consumer and give further power to the consumer such that they can innovate upon SDN technologies.

Second, despite the promise of a unique protocol as standard interface between the controller and the network infrastructure, interoperability between different controllers and network devices is hindered and closed ecosystems (being them based on open- or closed-source solutions) are emerging. Moreover, there is a need for a tighter integration between SDN controller platforms and cloud-based platforms.

Third, while current projects aim at providing a single integrated development environment to support the development lifecycle of SDN-based programs, more advanced solutions able to encompass other software domains like cloud and the Internet of Things in an holistic manner will be of utmost importance in the future to properly manage the complexity of such new generation of networked software-defined infrastructures.

Finally, it is believed that built-in security and efficient resource management should be taken thoroughly into account in future research projects.

4.4 Advanced NFV-based systems

Network Virtualization (NV) brings virtualization concepts to the network, similar to cloud computing, which was enabled by virtualization of servers. Network Functions Virtualization (NFV) focuses on virtualization of software-based network functions. Classical examples include virtualization of home gateways, firewalls, set top boxes, deep packet inspection components, IMS components, and monitoring probes. Instead of installing and managing dedicated hardware boxes for these functions, they are instead implemented as software components and deployed on commodity hardware infrastructures, in most cases operated by a network operator and referred to as telco clouds. Service Function Chaining (SFC) consists of building services using VNFs. Both network operators and service providers currently adopt the above principles at a sustained pace due to the many benefits brought by NFV (i.e. flexible, on-demand provisioning and rapid roll-out of telecom services with high re-use of available network and cloud resources).

Several areas remain open for interesting research. First, the exposure of cloud and network infrastructure for the deployment of virtualized network functions requires adequate infrastructure modeling frameworks as well as a mechanism to cope with available and reserved resources. Because the available infrastructure resource space is significantly larger than in the infrastructure of individual datacenters, or small networks, these mechanisms need to be extremely scalable, dynamic, and enable on-demand reservation mechanisms. In addition, complex modeling mechanisms are required in order to model the relationship of individual network functions and their potential applicability on available infrastructure. For example, suitable resource reservation models for interworking of private networks and public internet. This involves models for resource usage of network functions, as well as mechanisms to combine and isolate reservations on physical infrastructure, such that computing, memory and network capacity can be adequately reserved on the infrastructure.

Second, the definition of Service Function Chains requires a re-usable framework enabling to decompose services to network functions and service level agreements into constraints and characteristics of these network functions and their interactions. Currently there is no scientific foundation on how this can be achieved in a broader context. Mechanisms are needed to map particular

performance indicators to combinations of network functions and monitoring points in order to meet Service Level Agreements (SLA), which is considered to be commercially mandatory.

Third, protection or restoration at millisecond level of failed VNFs and servers using a 1:N scheme should be possible (i.e. no 100% duplication of all resources required), including protection of stateful VNF. The latter should be possible in multi-domain multi-operator scenarios as well.

Fourth, NFV-based networks are facing a rapid evolution toward complex software-driven systems. While most of the community is currently investing time and effort in proposing hybrid SDN/NFV architectures able to address as many network-specific requirements as possible, limited effort is devoted toward the deployment of tools to support SDN-/NFV-specific software development. Researching platforms to harmonize and integrate the development process of complex software-driven infrastructures is therefore highly desirable. This involves different platforms (e.g., debuggers and profilers), but also taking into account the multitude of different control protocols and controller platforms.

As an extension of the above, the full life cycle of composite network functions, services should be supported across both business and technical phases. This lifecycle should be automated as much as possible leading the way to quick and fast deployments and updates through the application of continuous integration mechanisms, widely used in the ICT domain. The orchestration should be resource and service aware, i.e. know which resource or service is optimum based on the technical needs of the application and the business criteria of the consumer. The orchestration process should be a continual process that seeks constantly to optimize the managed deployment. There are of course synergies with the focus of section 4.3.

Fifth, secure boot, configuration and operation of servers and VNFs at in-secure locations are considered to be very important for future research, together with hiding secure data in the VNFs/VMs from the hypervisor and server administrators.

Finally, the support of a VNF marketplace allowing the participation of third party developers fostering an open market environment, increasing the competition and incubating innovation presents a lot of challenges. Some of these challenges include automating the process of validating and certifying the provided VNFs, authorizing of the developers participating in the marketplace, providing auction mechanisms and elaborating on new SLA schemes and business models.

4.5 QoE-centric management

Europe is currently at the forefront of defining and measuring QoE in networked media applications, i.e. the quality as perceived by the end users. European companies have taken the lead in ITU.T standards development, and the COST action Qualinet has brought European researchers and industry together for larger dialogue. In spite of this leadership, we experience every day that proven and rational QoE metrics are set aside in favor of unproven or even disproven methods by research and industry.

It is commonly agreed by the ongoing FP7 Future Internet projects that future projects should focus more on optimizing QoE. For instance, a majority of faults reported to ISPs by their broadband

customers are due to issues in the home – for example low quality Wifi or home wiring or mis-configuration. There are currently very few existing measurements or tools to help diagnosing and mitigate the problems in an efficient way.

Automatic analysis of QoE measurements and automatic enforcement of actions based on observed context parameters (e.g., dynamic service adaptations) is put forward as an interesting topic for future projects. Measurement tools in current Future Internet projects produce interesting statistics about performance, both of the network and of the impact on a particular application. For instance, metrics and associated measurement methods to estimate the QoE of YouTube traffic are being designed. Today, the results are examined by hand, perhaps with the help of a dashboard (visualization tool). Future research projects should be dedicated to automatically identifying issues and isolating their cause. Especially when measurement capabilities are embedded in every home gateway and many end devices, there are interesting opportunities for automatic analysis and contextual action enforcement based on QoE measurements.

The development of objective QoE metrics and related QoE models generally applicable must be based on extensive subjective evaluations. Humans' perception of quality varies widely with technology, kind of application, context and content, which means that the required effort for developing metrics is magnificent, and on top of this, provides a moving target that requires updates as technical advancements appear. Overcoming these challenges in concerted, sustained Europe-wide actions, which unite the efforts of several relevant fields is a pre-condition for understanding the quality of future technical decisions. It is important to approach this challenge in a way that broad and affordable access is provided to proven QoE assessment tools for both research and industry, and stimulation of their ubiquitous adoption, is considered as very important for the success of future technologies.

Finally, the Future Internet Cluster believes that novel Device-to-Device (D2D) architectures will be NFV-based. As a consequence, the novel NFV-enabled D2D architecture will require flexible and intelligent management solutions for managing the provisioning of the service at an adequate quality level, creating the need to design and implement a QoE centric management framework, which will enable the optimization of network resources in terms of QoE, satisfying the different service and resource requirements of all the involved entities in the Device-to-Device (D2D) environments.

4.6 Advanced security for smart networks

As put forward by many of the Future Internet projects, security is addressed to some extent, but more detailed security related topics in smart networks should be addressed in future calls. In fact, many contributors are in favor of a few security-focused projects in the future. A commonly agreed topic, where dedicated research is interesting and required is: DDoS (Distributed Denial-of-Service) detection and prevention, which includes research on Firewall-as-a-Service (FaaS), DDoS protection networks, and DDoS information coordination between ISPs to jointly detect and protect DDoS attacks. Although the ideas and techniques behind DDoS attacks are known for many years, large-scale attacks (300 to 400 Gbps) did not happen until recently. The reasons behind the current explosion of DDoS attacks are the availability of easy to use attack tools and services (DDoS as a Service) and the fact that the potential of performing such attacks has been discovered recently by the broad community (including criminal

organizations). On the other hand, recent developments like SDN and OpenFlow provide novel approaches to mitigate DDoS attacks. Further research in this area is therefore needed.

Next to DDoS detection and prevention, also extrusion detection and prevention is also an important upcoming topic (i.e. detect when vital and private information is leaving networks triggered by malicious users, for instance in industrial espionage use cases, and design techniques to prevent this). Extrusion detection and prevention can be considered as the inverse use case of the well-known and studied intrusion detection and prevention use cases.

A third interesting topic is the ability to verify the actual neutrality and privacy preservation of ISPs, cloud providers and application providers, i.e. allow end-users to control that the providers do not modify, store and exchange the private data of Internet users. A measurement-based approach is most appropriate, which allows illuminating the currently obscure dynamics of the Internet by defining mechanism to identify policies violations, from simple SLA agreements, to privacy violations, to price discrimination, etc. Monitoring technologies are needed that allow third-party verification of the services available on the Internet, including the used networks and accessed applications. These technologies include active and passive monitoring system that can run specific tests to identify policy violations, by means of specific protocols that allow policies enforcement and verifications, where end-user terminals, networks, and cloud servers could be instrumented to cooperate.

4.7 Advanced Internet architectures

The current ongoing Future Internet projects on advanced Internet architectures focus on clean-slate approaches and address the fundamental limitations of the TCP/IP technology, reduce Internet latency and support an ever growing number of devices. The clean-slate initiatives provide solutions in a fundamental way, aiming at fixing deeper issues while reducing overall complexity, but the trade-off is they necessarily have to look at a much longer deployment horizon. The key challenge is to design a disruptive technology so that it does not disrupt the legacy technology it aims to replace upon deployment. The most important goal of a Future Internet Architecture is to provide a clean, scalable but at the same time simple structure that is able to support multiple internetworks with different characteristics tailored to different applications. Several interesting topics have been proposed by the Future Internet Cluster to be addressed in future research projects.

First, attention should be devoted to reducing end-to-end latency, by considering on the one hand clever queuing mechanisms and fast feedback from queues in network devices, and on the other hand by bringing both content and cloud-based applications much closer to the network edge for achieving low and predictable-latency networking. Latency is an end-to-end problem that cuts across all layers of the communication stack and most components of the communications infrastructure.

Second, it is believed that Future Internet architectures should be more anticipatory in nature, i.e. instead of reacting to changes, networks should anticipate to network changes and prepare the network elements to deal with changes. This topics has only be explored in some very initial ways: anticipation happens today on very short time scales in mobile radio networks (e.g., channel prediction) and is a tool

in green networking (e.g., to turn off base stations). Some very initial ideas exist in mobile video streaming or in data center networks. A thorough and systematic investigation of anticipatory networking will give a significantly better understanding of what kind, accuracy, and precision of anticipated information is needed, and to better design corresponding networking architectures.

Third, Future Internet architectures should devote specific attention to built-in resilience, i.e. react and anticipate to large or small scale damage caused due to intentional or unintentional acts. Damage to a smart network infrastructure should not prevent people and devices from communicating with each other. Furthermore, built-in security and advanced mobility support should be key design objectives in Future Internet architectures.

5. Conclusions

This position paper was discussed and documented by the ongoing FP7 Future Internet projects and the followed methodology has been described in this paper. Based on many interactions and discussions, the following important topics for future research (not covered in ongoing projects, but considered very important for future research and commercially relevant in turn) have been identified and reported upon in this position paper: Advanced content delivery systems, measurement-based management, SDN-based systems and applications, advanced NFV-based systems, QoE centric network management, advanced security for smart networks, and advanced Internet architectures.

We are very grateful to all contributors for their timely and inspiring inputs.

For feedback or more information, please contact the editors: Prof. Filip De Turck (*filip.deturck@intec.ugent.be*), Prof. Thomas Michael Bohnert (*thomas.bohnert@zhaw.ch*), and Dr. Antonio Cimmino (*cimm@zhaw.ch*).

6. References

[1] FP7 BUTLER project, http://www.iot-butler.eu/, Okt. 2011 – Okt. 2014

[2] FP7 CONCERTO project, http://ict-concerto.eu/, Dec. 2011 – Nov. 2014

[3] FP7 eCOUSIN project, http://www.ict-ecousin.eu/, Nov. 2012 – Apr. 2015

[4] FP7 FLAMINGO project, http://www.fp7-flamingo.eu/, Nov. 2012 – Oct. 2016

[5] FP7 FUSION project, http://www.fusion-project.eu/, Jan. 2013 - Dec. 2015

[6] FP7 GreenICN project, http://www.greenicn.org/, Apr. 2013 – Mar. 2016

[7] FP7 LEONE project, http://www.leone-project.eu/, Nov. 2012 – Apr. 2015

[8] FP7 MCN project, https://www.mobile-cloud-networking.eu/, Nov. 2012 – Oct. 2015

[9] FP7 mPlane project, http://www.ict-mplane.eu/, Nov. 2012 – Okt. 2015

[10] FP7 NetIDE project, http://www.netide.eu/, Jan. 2014 – Dec. 2016

[11] FP7 ONE project, http://www.ict-one.eu/, Sept. 2010 – Aug. 2013

[12] FP7 PACE project, http://www.ict-pace.net/, Nov. 2013 – Okt. 2015

[13] FP7 PRISTINE project, http://ict-pristine.eu/, Jan. 2014 – Jun. 2016

[14] FP7 RITE project, http://riteproject.eu/, Nov. 2012 – Nov. 2015

[15] FP7 ROMEO project, http://www.ict-romeo.eu/, Okt. 2011 - Sept. 2014

[16] FP7 SmartenIT project, http://www.smartenit.eu/, Nov. 2012 – Oct. 2015

[17] FP7 T-NOVA project, http://www.t-nova.eu/, Apr. 2014 – Dec. 2016

[18] FP7 Trilogy 2 project, http://trilogy2.it.uc3m.es/, Jan. 2013 – Dec. 2015

[19] FP7 UNIFY project, https://www.fp7-unify.eu/, Nov. 2013 – Apr. 2016