SUPERFLUIDITY

A SUPER-FLUID, CLOUD-NATIVE, CONVERGED EDGE SYSTEM

Research and Innovation Action GA 671566

DELIVERABLE D8.3:

INNOVATION AND EXPLOITATION PLAN

Deliverable Type:	Report
Dissemination Level:	Public
Contractual Date of Delivery to the EU:	01/07/2016
Actual Date of Delivery to the EU:	02/07/2016
Workpackage Contributing to the Deliverable:	WP8

Editor(s):	George Tsolis (CITRIX), Bessem Sayadi (NOKIABLF), Carlos Parada (ALB), Raúl Álvarez Pinilla (TELCA), Rufael Mekuria (USTR), Philip Eardley (BT), Michael J. McGrath (INTEL), Maria Bianco (CNIT), Nicola Blefari Melazzi (CNIT), Erez Biton (NOKIAIL), John Thomson (ONAPP), Christos Tselios (CITRIX)
Author(s):	George Tsolis (CITRIX)
Internal Reviewer(s)	Nicola Blefari Melazzi (CNIT)

Abstract: This deliverable holds the plan on identifying and driving the innovation and commercial potential for the technical work and technical goals of the project, providing common framework for innovation and driving towards commercial exploitation for the efforts and results from all technical work packages. Moreover, the plan describes the process by which identification of potential innovation and commercial exploitation inhibitors, within and between the work packages, is taking place.

Keyword List: Innovation, Exploitation

Version Control Table			
Version N.	Purpose/Changes	Author	Date
1	First release	SUPERFLUIDITY Project	01/07/2016
1.1	Revision	Nicola Blefari Melazzi	01/07/2016

INDEX

List of Tables

Glossary

SUPERFLUIDITY DICTIONNARY	
TERM	DEFINITION
Innovation	The act or process of introducing new ideas, technologies, or methods.
Exploitation	Direct or indirect utilisation of foreground in further research activities other than those covered by the project, or for developing, creating and marketing a product or process, or for creating and providing a service, or for using them in standardisation activities. (Article 28.1 of GA)
Results	Any tangible or intangible output of the Action, such as data, knowledge and information whatever their form or nature, whether or not they can be protected, which are generated in the Action as well as any rights attached to them, including Intellectual Property Rights.
General Public	Audience composed by people, groups and organizations which do not have specific competences/knowledge in the matter of the project but only a general interest at the application of results in society at a large.

Table 1: SUPERFLUIDITY Dictionary.

1 Introduction

1.1 Deliverable Rationale

This deliverable documents the activities of Task 8.3 (T8.3) of Work Package 8 (WP8), which specifies:

Work Package 8: Communication, Dissemination, Standardization and Exploitation

Task 8.3: Innovation and Commercial Exploitation

This task will be in charge of Innovation and Commercial Exploitation management. It will include the following activities:

- Overall project responsibility for identifying and driving the innovation and commercial potential for the technical work and goals of the project.
- Ensure of a common approach for Innovation and Commercial Exploitation in all work packages while maintaining the overall vision of the Project.
- Identification of potential innovation and commercial exploitation inhibitors in and between the work packages while implementing the technical work.
- Preparation of proposals for the Management Board on innovation concepts.
- Preparation of summaries for the periodic reports and annual reports, and for the reviews.
- Delivering of relevant presentations both internally and externally to the industry.

1.2 Executive summary

1.2.1 Deliverable description

This deliverable describes the plan for identifying and driving the innovation and commercial potential for the technical work and technical goals of the project, providing common framework for innovation and driving towards commercial exploitation for the efforts and results from all technical work packages. Moreover, the plan describes the process by which identification of potential innovation and commercial exploitation inhibitors, within and between the work packages, is taking place.

1.2.2 Deliverable structure

Section 2 covers Innovation Management. More specifically, Section 2.1 outlines the management structure, and how it fosters innovation through enabling collaboration between the partners. Section 2.2 outlines the state of the art, per our review and continuous monitoring. Section 2.3 outlines the innovation pathways the projects envisions to take, as well as early results. Last but not least, Section 2.4 summarises the innovation potential of the project.

Section 3 covers Exploitation Planning. Section 3.1 outlines the exploitation plans of the individual partners, while Section 3.2 summarises the exploitation plan of the consortium as a whole. Finally, Section 4 sets expectations in terms of planned deliverables on these matters. The deliverable concludes with a list of References, related to both the state of the art and the project results so far.

2 Innovation Management

2.1 Introduction

Activities in this area aim to fulfil the Scientific and Technological Vision of the SUPERFLUIDITY project (please refer to section 3.2 of deliverable *D1.2 Project Vision and Roadmap*).

To foster innovation, the project has to, first of all, implement the appropriate management structure and provisions for frictionless collaboration between the partners. Subsequent paragraphs of this section provide more information.

Secondly, the project has to keep pace with the ongoing scientific research relevant to the technological objectives. Section 2.2 outlines the state of the art, per our review and continuous monitoring, organised in the following areas:

- Cloud Networking
    - Cloud-based, Virtualised Network Services
    - Software Radio and DSP Virtualization
    - High Performance, Software-based Packet Processing
- Network Services Decomposition and Programmability
- Cloud RAN and Mobile Edge Computing
- Automated Security and Correctness

Thirdly, the project has to advance the state of the art, in research directions that are aligned with the project scientific and technological vision. Section 2.3 outlines the innovation pathways the projects envisions to take. We also leverage the opportunity to report early results on this front.

Finally, and as an extension to the research directions and innovation pathways covered above, Section 2.4 summarises the innovation potential of the project. Most importantly, this is associated to specific market segments of interest to the European CSPs, industrial partners and SMEs of the project, where the innovations of SUPEFLUIDITY are expected to result in competitive advantages.

2.1.1 The Management Structure

As further elaborated in deliverable *D1.1 Project Management Manual*, the Management structure has been conceived in such a way as to ensure flexibility and dynamism among project bodies in order to avoid any stiffening effect in technical performance. On the other side it guarantees a strict control of operational and administrative/financial compliance.

The skeleton of the structure is made up of two boards:

- General Board: highest decision body of the project, taking formal and legally binding decisions, representing the interests of the project partners.
- Management Board: in charge of the executive management of the project.

Following innovation management principles [ARN2010, MAI2012], the structure is completed with additional bodies, which cooperate with the Management Board, namely:

- Scientific Team: inspires the project research vision, its main challenges, approaches and objectives, updating them in time. It gathers ideas from project researchers, elaborates and circulates them among all WP members.
- Business and Marketing Team: brings into the project requirements from business units, validates project ideas from a business potential point of view; streamlines project work towards exploitable results; exports project results within partners' business units as a first step toward real markets.
- Review Team: reviews all project outputs (mainly deliverables) but in general all important documents and results released outside the project. It is responsible for overall quality control.

All in all, we believe that the above structure keeps the advantage of usual EU projects practises, as we have known them; is more adherent to reality; brings innovation at the centre of the process; acknowledges demands and inputs from business units and from advisors; takes results to the outside world, always with impact in mind; reviews project products and deliverables to assure their quality, while at the same time complying with contractual agreements.

2.1.2 The Innovation Coordinator

The Innovation Coordinator assists and advises the project in how to best respond to emerging market opportunities.

The Innovation Coordinator:

- has the overall project responsibility for responding to emerging market opportunities, assuring that the project research ideas and achieved scientific results can be transformed into products and services;
- prepares proposals for the MB on exploitation-related issues;
- prepares exploitation summaries for the periodic (four-monthly) reports and annual reports, and for the reviews;
- gives presentations both internally and externally.

The Innovation Coordinator legal entity is CITRIX. The Innovation Coordinator person is George Tsolis, a Principal Architect with fifteen+ years of business experience in appliance software, mobile telecommunications and cloud industry.

2.2 State of the Art

In this section we cover a set of research areas related to SUPERFLUIDITY, mention the current state of the art and comment where the current issues in each lie.

2.2.1 Cloud Networking

2.2.1.1 Cloud-based, Virtualised Network Services

Cloud computing has moved beyond a buzzword to being an integral part of everyday life, with most online services being provided either from a private data center or a public cloud. Network virtualization is typically implemented in clouds by providing virtual network interface cards (NICs) to customers that are directly mapped to a VLAN by an hypervisor (e.g., NICIRA/VMWare), thus ensuring traffic isolation by default. Network virtualisation is also available in the Internet, with MPLS being the protocol of choice. A number of network providers and the larger cloud providers also provide virtual network overlays including Amazon (http://aws.amazon.com/vpc/), but these are normally used as VPN-type systems to segregate networks and join particular resources together. Some vendors are working just in this space, including PlumGrid (http://plumgrid.com/) who specialise on Virtual Network Infrastructure.

The advent of OpenFlow/SDN offers another alternative, promising to create dynamic network services on the fly, both over the local area and over the wide area between data centres and cloud providers. Juniper's implementation of OpenContrail as a Juniper Contrail is a network virtualisation platform using software defined networking. A key characteristic of public clouds is that processing is localised to one of a few locations for each cloud provider; for instance, Amazon allows clients to currently choose between US east and west coasts, and Western Europe.

Even more recently, network services and packet processing have started acquiring a distinctly cloud flavour. Proof of this is the rise of Network Function Virtualization (NFV [NFV2014]), an approach heavily backed by major world-wide operators and service providers which aims to develop the standard IT virtualization technology found in clouds and data centers in order to implement network services (anything ranging from middleboxes such as firewalls, NATs, IDSes and scrubbers to more service-oriented boxes such as content caches). Even more holistic approaches exist such as ECOMP [ATT2016], a framework for real-time, policy-driven software automation of network management functions, designed and build for carrier scale workloads. ECOMP is currently optimised to deliver essential management tools for operating services upon virtualised network and infrastructure resources, facilitating rapid introduction of computing, memory and storage resources, used to dynamically instantiate, operate and lifecycle manage services, virtualised network functions and smart cloud infrastructure. Being an open source platform, it is capable of supporting any business domain and generates value by enabling automated network function virtualization and by supporting the interoperability across third party cloud infrastructures and networks. However, these are early days for NFV, and major hurdles remain.

First, hardware based middleboxes, despite their shortcomings, are good at achieving high, carrier grade (10-40Gb/s) performance. However, while recent advances in software-based processing on commodity hardware have made significant progress towards achieving 10Gb/s rates for most packet sizes, pushing that envelope further in virtualised environments, while ensuring low packet drop rates remains an open problem. More recently, research has looked into the viability of running virtualised middleboxes in the cloud [SHE2012b] by showing that the "detour" of deploying such functionality at a somewhat remote data center would not overly affect end-to-end delays. Further, Comb [SEK2012] introduces an architecture for middlebox deployments targeted at consolidation. However, it does not support multi-tenancy nor isolation, and the performance figures reported (about 4.5Gb/s for two CPU cores assuming maximum-sized packets) are far from carrier grade. Finally, ClickOS [MAR2014], developed by NEC, looked at combining lightweight host virtualization with Click to provide high performance and scalability for network function virtualization. Other work [MAD2013] (as well as Erlang on Xen, LuaJIT and HalVM) has similarly focused on minimalistic virtual machines, but did not specifically target network functionality nor high performance. The overall traction that NFV has gained [MIJ2016a], led to its adoption by numerous vendors and marker leaders [VZN2016], which proceed towards implementing NFV-, related solutions in their frameworks, mostly based on the directions described in [NFV001]. Recently, OpenStack Foundation endorsed NFV claiming that NFV and OpenStack can provide an agile, scalable, and rapidly maturing platform with compelling technical and business benefits for telecommunications providers and enterprise network operators [OSF2016], while a scalable and application-agnostic scheduling framework for packet processing, directly related to the NFV principles is described in [PAL2015].

2.2.1.2 Software Radio and DSP Virtualization

Wireless access virtualization enables diversity and co-existence of different wireless technologies on the same hardware, which offers the network providers/owners the advantage of simultaneous operation of multiple standards, multiple base stations and multiple tenants on same platform [YAP10] [SACH08]. Depending on the performance and sharing capabilities of underlying hardware, various virtualization approaches regarding protocol stack (L1-L3) have been recently proposed. Inspired by the SDN concept, the most common is the flow-oriented approach often referred to as mobile network virtualization, in which the physical downlink and uplink blocks are abstracted as network elements, and a virtualization layer attempts to manage the data flows associated with these elements. Examples are OpenRoads [YAP10] and vBTS [BHAN10], where virtualization is implemented as a software switch overlay, however, enabling the control of radio resource scheduling for service level guarantee is a problem, which can be solved by exploiting a wireless hardware scheduler as proposed in NVS [KOK10] and Virtual LTE [ZAK10]. Unlike the flow-oriented approach, the protocol-oriented approach attempts to instantiate multiple versions of a protocol stack on the same hardware platform. Software and/or hardware resources (e.g., CPUs and DSPs) have to be sliced in this case as proposed in [SACH08]. Partial implementations supporting single protocol virtualization by protocol enhancements have been proposed in [KOK10] [ZAK10]. In contrast, full implementations (i.e., virtualization of multiple different protocol stacks) are only possible using SDR technology as introduced in OpenRadio [BAN12] and Sora [TAN11]. Overall frameworks such as Atomix [BAN15] are also introduced, demonstrating that it is feasible to build modular DSP software by building the application entirely out of fixed-timing computations, an approach that alleviates programmers towards deploying new signal processing applications. Beside flow- and protocol-oriented virtualization, RF-level and spectrum-oriented approaches have been proposed, which are based on sharing single RF circuit by multiple baseband protocols and by exploiting cognitive radio concepts for spectrum usage. This concept, however, requires modifications of RF circuits in order to enable radio hardware virtualization. Moreover, radio waveform modifications for spectrum sharing are not supported in current and on-going standards and, for that reason, these techniques are beyond the scope of this project.

2.2.1.3 High Performance, Software-based Packet Processing

Networking equipment has traditionally been implemented using ASICs due to stringent performance requirements. However, this results in inflexible, as well as proprietary, embedded systems. This, coupled with the need for agile network functionality deployment and the availability of high performance commodity hardware (e.g., x86 servers with 1-10Gb NICs) in data centers, have led to a wealth of recent work on software-based packet processing functions. Netmap [RIZ2012,RIZ2012a], and DPDK [DPDK2016], developed by INTEL, are examples of very high performance host networking frameworks, designed to provide high-speed, low overhead user-space I/O packet channels, and as such, are poised to form the basis of many software-based packet processing systems. Indeed, such frameworks enable the use of custom network stacks, specialised for the task at hand, resulting in increased application performance compared with the use of generic network stacks as is the case in today's operating system kernels [MWH2014].

Several platforms have also been proposed to provide routing functionality on commodity servers. Click [KOH2000] is a modular system to quickly build pure software network processing pipelines, providing high flexibility at the expense of reduced performance. RouteBricks [DOB2009] proposed a high performance distributed routing platform based on clustering. PacketShader [HAN2010] was the first routing platform that proposed the joint use of CPUs and GPUs for networking systems. While PacketShader was thus a heterogeneous network system, it only targeted those two specific components, resulting in limited heterogeneity. Last but not least, NetVM [HWA2015] brings virtualization to the Network by enabling high bandwidth network functions to operate at near line speed, while taking advantage of the flexibility and customization of low cost commodity servers.

While all this work on networking for commodity hardware has shown that such systems can provide some networking functionality, a sharp decrease in performance has been observed as the complexity of the network processing (and its virtualization) increases: either the aggregate performance (throughput) decreases, the aggregate memory footprint grows, or some other performance metric, such as power-per-bandwidth, quickly becomes untenable (e.g., in clusters-based systems such as RouteBricks).

2.2.2 Network Services Decomposition and Programmability

Defining the right set of building blocks is a necessary starting point towards providing a truly virtualised, integrated and transparent 5G core and access network infrastructure. Most of the NFV related works focuses on the strictly engineering and infrastructural challenges of (1) how NICs and switches can be virtualised for seamless migration to different VMs [PICSIG, 8021Q, RFC6325]; (2) how virtual LANs deployment can provide customer traffic isolations and geographically distributed LAN composition [NVGRE, VXLAN, STT, NVO]; (3) virtual network I/O performance optimization [RAM2009, DON2012, DPDK2016] (4) energy consumption optimization [BEL2010, BOT2012] and (5) primitive composition languages and platforms. Similarly, most of the radio function virtualization works focus on implementation and composition [or on more technical details like efficient time synchronization or efficient demodulation [MUE2010, WU2013, BHA2013, GUD2013].

As a starting point, we could take the ETSI NFV specification “Methodology to describe Interfaces and Abstractions” [NFV007], which makes use of functional blocks (and their interconnection) as system composition abstraction (the Click modular router is another example of this [MOR99]). However, some reasonable questions may rise: is the functional block abstraction the most suitable? Do we need to further partition a single functional block? How do we define the sub-procedures that the functional blocks consist of?

At least three possible approaches for a more granular network functions definition exist. First, the most intuitive one is the match action rule set, which is adopted by Openflow compliant data planes (among other software-defined approaches, like [BOS13, JEY13, SON13]). Openflow has so far undergone many revisions, and the latest version [OF1.4] extends such an elementary abstraction with a large number of apparently ad hoc extensions such as more flexible header matching, action bundles, multiple pipelined tables, synchronised tables, support for multiple controllers, etc. (with newer extensions such as typed tables [FAWG] and flow states [CRA13] currently under discussion). Limited attention has so far been posed on to how to rethink the original programmatic abstraction (besides the abandoned Google OpenFlow 2.0 proposal [YAD11], considered too disruptive [Mey13]) so as to accommodate such extensions in a clean API [PER13]. While OpenFlow’s abstraction results perfect for many applications like Ethernet and IP forwarding mechanisms, complex, more dynamic and real-time applications may require a somehow more versatile programming abstraction. For this reason, more recent works in the fields of wireless MAC programming [TIN12, BIA12], programmable monitoring platforms [BiaArx14] and software defined networking [BIA14, MOS14] have proposed stateful approaches based on state machines to describe complex networking related complex tasks like wireless MAC mechanisms, cyber-threat detection/mitigation and layer 2/3 protocols. Complex middlebox applications still need more flexible programming approaches and therefor solutions based on application specific language [MON13, FOS11] or Application Programming Languages (API) bound to generic purpose programming languages may be needed. Solutions like the C API provided by Open Dataplane [ODPL] and PacketC [DUN09] in facts represents an attempt to provide a high-level language to simplify the construction of network processor-based applications.

2.2.3 Cloud RAN and Mobile Edge Computing

A generic MAC architecture for scheduled systems was suggested in FLAVIA [MAE11]. The suggested architecture relies on modules, each representing a service such as admission control and load balancing, that a scheduled access technology is expected to provide. However, while it tackles the problem of decomposition, it forces a standard hardware platform for each access point which is wasteful ,since every function primitive is implemented on each access node (e.g., eNodeB). In [GUD2013], Gudipati et al. propose SoftRAN, a software-defined centralised control plane for radio access networks. SoftRAN adopts a 2-tier model where part of the control stays within the nodes (the part that requires frequent decisions with a local scope) and part goes to the central controller (the part that requires less frequent decisions but with a more global scope). In contrast, SUPERFLUIDITY proposes a more fluid approach where part of the RAN services can be dynamically migrated from the node to the cloud and vice-versa. The management and orchestration of those services is converged with the cloud infrastructure management system, leveraging the quick instantiation and performance enhancements of SUPERFLUIDITY. In [PEN13], Pentikousis et al. advance the state of the art in SDN by introducing software-defined mobile network (SDMN) architecture, where they employ network virtualization. The main theme of [PEN13] is the decoupling of data and control planes relying on SDN which fosters the capability of the operator to define new services without the need to replace hardware or to go through a tedious implementation trail. It appears that C-RAN will enable energy efficient network operation and possible cost savings on baseband resources. Furthermore, it improves network capacity by performing load balancing and cooperative processing of signals originating from several base stations [CHE15]. This renders the particular architecture really appealing to both industry and academia, as shown in [HUA14], while numerous white papers push C-RAN integration constantly further [NGMN15,RAD15]. Much more recently, a set of major international operators and vendors (INTEL is one of them) has started the ETSI ISG on Mobile Edge Computing [MEC], which advocates for the deployment of virtualised network services at remote access networks, next to base stations and aggregation points, running on x86 commodity servers. These are early days for this proposal, but their goals are well aligned with those of SUPERFLUIDITY, and so the ISG gives the project an interesting standards venue for disseminating its technical results.

A somehow similar approach towards converging mobile and cloud computing is presented by [OEC], [SAT15] with the introduction of cloudlet, a new architectural element which can be viewed as a "data center in a box". The notion of cloudlet is based on four key attributes (i) is soft state dependent (ii) remains always powerful, well-connected and safe (iii) obtains logical proximity to end devices, resulting into low end-to-end latency and high bandwidth (iv) builds on standard cloud technology. These attributes allow cloudlets to operate in a self-managing fashion, due to the increased caching capabilities, possess sufficient compute power to offload resource intensive computations and incorporate all classic cloud infrastructure deployment functions, thus being compatible with leading cloud infrastructure/frameworks such as Amazon EC2 and OpenStack. The proposed OpenEdgeComputing architecture bears many similarities to the original notion behind SUPERFLUIDITY platform design, however lacks the flexibility and function re-usability that SUPERFLUIDITY envisions.

2.2.4 Automated Security and Correctness

When designing new virtualised environments, there are two main goals in terms of privacy and security: first, tenants of the virtualised environment must be isolated from each other, both from a security as well as from the performance point of view. Second, tenants must be prevented from being able to harm the cloud or the Internet at large. To achieve these goals, public clouds use virtualization (typically Xen, KVM or Hyper-V), which provides CPU/memory and disk security isolation and (a certain level of) performance isolation.

Cloud *network* security has two facets: access control and performance isolation. Access control lists can become quite complex, as they have to capture a myriad of policies (tenant-Internet, tenant-tenant and tenant-cloud services communication). Traditionally, enforcing such ACLs is done using a mix of firewall rules applied to tenant traffic (in dom0) and VLANs, yet it has been argued that these do not scale very well to large number of tenants and the very high virtual-machine churn seen in data centers today. An alternative, proposed by CloudPolice [POP2010] is to enforce fine-grained security rules in the hypervisor on-demand, only when two virtual machines want to communicate. Additionally, shaping is used to ensure that tenants do not affect each other's performance. For instance, in Amazon EC2 high-bandwidth traffic from lower-cost VMs is shaped to a seemingly fixed allowance [RAI2012, LAC2013]. The downside of such simple static shaping is wasted capacity, and a number of research papers have discussed alternative sharing models including FairCloud [POP2012], EyeQ [JEY2013] and ElasticSwitch [POP2013]. An architecture where software-based, virtualised services can be instantiated whenever, wherever they are need, as proposed by SUPERFLUIDITY, would present a number of **new security issues** that stem from the **interactions of the virtualised services with operator-deployed functionality**. Checking whether operator security policies are obeyed can be done at runtime or at install time, but **runtime enforcing imposes high costs**. We propose to advance the state-of-the-art in checking and enforcing security policies by suggesting novel detection tools which will allow real time evaluation of tenants behaviour, whether compared to a given policy to be enforced or, in more involved circumstances, compared to a learned normal behaviour of the system, that is, identify abnormal behaviour without the strict a-priori definition of a normal behaviour and without any prior knowledge of how such a behaviour is characterised.

Checking for compliance (both in terms of security and correctness) at runtime is typically done by injecting packets and observing the outcome. Ping and traceroute are the default tools of administrators, but they will not scale for SUPERFLUIDITY. Automatic generation of test packets is scalable as long as only static forwarding rules are being checked [ZEN2012]; this work does not address the issue of stateful middleboxes, prevalent today in operator networks, and soon to be ubiquitous due to in-network clouds. Probing techniques have been developed to test the behaviour of middleboxes when presented with certain flavours of TCP traffic in the TCPExposure tool by NEC and UPB [HON2011] that runs at both the client and the server. TraceBox [DET2013] is another tool that uses a traceroute-like approach to detect middleboxes with tests run only at the client; TraceBox is easier to deploy but is less accurate than TCPExposure. Simple [QAZ2013], goes a step further and tries to infer what a middlebox is doing by treating it as a black box. Relying entirely on checking would be foolish, and so SUPERFLUIDITY will further include a novel method to both learn behaviour as well as efficiently compare incoming traffic to the model learned. Both learning and testing of new traffic will be based on simple, easy to collect features (such as timing data, sizes, etc.) that will not require costly DPI or complex tests of pre-defined policies. Moreover, the baseline system will scale in several dimensions, monitoring, tracking and alerting not only with relation to tenants' communications, but also with relation to the way they use other system resources such as computation or memory.

2.3 Progress Beyond the State of the Art

In this section we cover what the project's progress beyond the state of the art is envisioned to be, as summarised in Table 2. Also, we report early results that demonstrate progress to that direction.

Research Area		SUPERFLUIDITY Goals
Cloud Networking	NFV Virtualised Environment	- Achieve 10-40 Gbps throughput with virtualised and software based packet processing. - Achieve tens of ms. (or lower) service instantiation. - Achieve massive consolidation and migration of services.
	Software Radio and DSP Virtualization	- Unify and abstract the protocol stack and the hardware platform to instantiate multiple vendor protocols. - Open radio platform and enable third parties to design the protocol stack. - Design portable runtime dataflow engines, enabling simultaneous seamless execution of multiple protocols.
	High Perf. Software-Based Packet Processing	- Achieve high performance with commodity hardware. - Increase the level of flexibility by going beyond the paradigms of proprietary, embedded systems. - Enable agile network function deployment.
Network Service Decomposition and Programmability		- Introduce program abstractions specifically targeted to 5G functions. - Combine block-based composition abstractions (such as those exploited in Click routers [KOH2000], or emerging in the ETSI NFV work on service chaining) with event-driven programming paradigms such as basic match/action based approaches or more powerful stateful abstractions based on extended finite state machines.
Cloud RAN and Mobile Edge Computing		- Enable modular "hot" replacing of eNB functions (such as scheduling) and allow migration of such functions between edge clouds and the antenna subsystem, so as to balance algorithmic complexity with front-haul capacity. - Enable the migration of non-RAN functions (like local caching and CDN) between the Remote Radio Head and the edge cloud, to maximise their performance.
Automated Security and Correctness		- Provide a pre-deployment checking system to ensure that virtualised network services do not negatively affect the network nor other tenants; the system has to be both scalable and stateful, able to model most types of services. - Implement a post-deployment system that will learn the behaviour of traffic and detect any anomalies, thus providing a further security

	mechanism in cases where the checking system does not have information about the processing performed by a network function, or when static analysis is inaccurate.

Table 2: SUPERFLUIDITY Innovations and Goals

2.3.1 Cloud Networking

2.3.1.1 Cloud-based, Virtualised Network Services

SUPERFLUIDITY will push the boundaries of the previous work with the goal of developing mechanisms in support for a superfluid, virtualised network architecture. Specific goals include, but are not limited to, achieving high throughput (10-40Gb/s and higher throughput), near-instantaneous virtualised service instantiation and migration (in tens of milliseconds or less) and massive consolidation of services (e.g., thousands of virtual machines on a single x86 server).

2.3.1.2 Software Radio and DSP Virtualization

In contrast to previous work, SUPERFLUIDITY addresses the question of how to unify and abstract both the protocol stack as well as the hardware platform in order to allow (dynamic) instantiation of multiple vendor protocols. SUPERFLUIDITY attempts to open the radio hardware platform and enable third parties to design protocol stacks by defining a set of primitives and associated APIs enabling protocol modelling and construction according to dataflow paradigm. In addition to this, SUPERFLUIDITY will design the portable runtime dataflow engine enabling simultaneous seamless execution of multiple protocols. In order to cope with computational and real-time requirements of protocol stack processing, hardware acceleration and its abstraction will complement the work in this area.

2.3.1.3 High Performance, Software-based Packet Processing

SUPERFLUIDITY proposes to advance the state of this art by considering the use of highly heterogeneous hardware for network processing, to open up levels of performance and power consumption tuning that could only be seen in ASICs, while retaining the flexibility only afforded by software, a proposition particularly attractive for ISP clouds. This includes leveraging hardware just now hitting the market such as low-power, small-sized microservers with architectures beyond the traditional x86 (e.g., MIPS64, ARM, ARM's BIT.little, to name but a few) in order to severely reduce energy costs, especially for edge deployments. Network and function Management and Orchestration (MANO) is of paramount importance for SUPERFLUIDITY, therefore we aim to enhance existing, yet still basic orchestration platforms [MIJ2016b].

The ability to automatically extract high performance out of commodity hardware remains largely an open research problem, and one that is only likely to receive greater attention as an increasing amount of network functionality moves to software.

SUPERFLUIDITY will aim at meeting the stringent requirements imposed by future 5G networks by designing and implementing a superfluid, converged network architecture that is location, hardware and time-independent. The work will push the boundaries of what is currently possible with virtualised, software-based packet processing (10-40Gb/s and higher, extremely fast service instantiation and migration in milliseconds, massive numbers of concurrent virtualised services on a single platform, significant power reductions, etc.). The goal is to bring the advantages of cloud and software-based networking to 5G networks so that services can be deployed whenever and wherever they are needed, and to leverage the availability of inexpensive, off-the-shelf hardware in the process. In addition, SUPERFLUIDITY will contribute to the enhancement of underlying application delivery methods through novel mechanisms, thus facilitating seamless end-to-end user experience [DET2016, MEK2016a, MEK2016b].

2.3.2 Network Services Decomposition and Programmability

SUPERFLUIDITY will devise programming abstractions specifically targeted to 5G functions. The API design work will address three programming levels: service, function, and processing levels, and will attempt to maximise viability by reusing existing standard (when applicable) or research community best practices. Work will on one side target the definition of 5G specific actions and events, and on the other side will address the specification of the constructs needed to combine and orchestrate a desired execution of such actions (conditioned on the arrival of events). Particularly promising and forward-looking is SUPERFLUIDITY's approach of combining block-based composition abstractions (such as those exploited in Click routers, in some software defined radio architectures, or emerging in the ETSI NFV work on service chaining) with event-driven programming paradigms such as basic match/action based approaches or more powerful stateful abstractions based on extended finite state machines [ETT2016].

2.3.3 Cloud RAN and Mobile Edge Computing

Beyond the current vision of a static RAN function fully located in one "edge computing" place, SUPERFLUIDITY will support the ability to modularly "hot" replace eNB functions (such as scheduling) and to permit migration of such functions between edge clouds and the antenna subsystem, so as to balance algorithmic complexity with fronthaul capacity. SUPERFLUIDITY will also transcend current Mobile Edge Computing vision where non-RAN functions (local caching, CDN,...) are envisaged to be co-located only at the eNB by enabling their migration between the RRH and the edge cloud, to maximise their performance. In addition, SUPERFLUIDITY will exploit the overall Mobile Edge

Computing architecture to design, implement, deploy and evaluate additional functional blocks for monitoring purposes [TSE2016].

2.3.4 Automated Security and Correctness

SUPERFLUIDITY will provide a two-pronged, complementary approach to security. First, it will go beyond the state of the art, providing a pre-deployment checking system that will ensure that virtualised network services do not negatively affect the network nor other tenants; unlike approaches in the literature, the system will be both scalable and stateful, able to model most types of services. Second, SUPERFLUIDITY will implement a post-deployment system that will learn the behaviour of traffic and detect any anomalies, thus providing a further security mechanism in cases where the checking system does not have information about the processing performed by a network function, or when static analysis is inaccurate.

2.4 Innovation Potential

SUPERFLUIDITY's ambition is to aid mobile network operators to meet the demanding requirements of 5G while at the same time develop favourable market conditions which will create sustainable business for all players in the value chain. In an environment of slowing market expansion rates and declining revenues, where end users request personalised services, better performance, better user experience, businesses need to get more information about their consumers, easier and secured access to devices and greater flexibility for provisioning new services. There is a key role to play for Equipment providers, Service Providers and IT players together to make this a reality by providing **converged IT and Network infrastructure**.

SUPERFLUIDITY can provide this, thanks to the unique blend of Telco and IT players in its consortium.

In this context, SUPERFLUIDITY helps to create a new ecosystem and the opportunity for all players to collaborate and develop new business models they can each benefit from. More in detail, SUPERFLUIDITY innovation is to:

- Enable Mobile Network Operators (MNOs) to rapidly deploy new services for consumer and enterprise business segments which can help them differentiate their service portfolio and to adapt them instantaneously (**time-independence**) to varying traffic demands by localising processing (**location-independence**), providing users with the best quality of experience constantly;
- Boost the capabilities of software and application providers by allowing them to develop applications that can take advantage of the information (through specific APIs) on radio network capabilities and conditions available at the base station without the need to worry about hardware (**hardware-independence**), performance and scaling issue (**scale-independence**).

Regarding specific market segments, the main innovation arising from a careful analysis of SUPERFLUIDITY innovation potential is that nobody in the market so far can claim to have even roadmapped such a vision (convergence of IT and Network infrastructure all the way from the core towards the access, open APIs for exploiting radio network information, near-instantaneous relocation of services, independence from hardware components, etc.).

Below we report an extract of such an analysis to substantiate our statements:

- “LTE Network Infrastructure”
    - **Leaders** (per Gartner, April 2015): Ericsson, Nokia, Alcatel-Lucent, Huawei (NEC is in the **Visionaries**)
    - **SUPERFLUIDITY innovation**: the converged architecture concept of SUPERFLUIDITY will guide network equipment vendors to evolve their infrastructure products to 5G.

- “Small Cell Equipment”
    - **Leaders** (per Gartner, October 2015): Huawei, Alcatel-Lucent, Nokia, Ericsson, Cisco, ZTE (NEC is in the **Challengers**)
    - **SUPERFLUIDITY innovation**: small cells will play a very important role in fulfilling 5G requirements. The advances in CRAN and MEC, as also contributed by SUPERFLUIDITY, will evolve these products through the adoption of softwarisation and cloudification.

- “x86 Server Virtualization Infrastructure”
    - **Leaders** (per Gartner, July 2015): VMware, Microsoft (Red Hat and Citrix are in the **Niche Players**)
    - **SUPERFLUIDITY innovation**: even though there are no plans to improve hypervisors, the adoption of minified VMs, unikernels, containers & specialised operating systems, and the reduction of orchestration system overhead, will bring those to the next level.

- “IT Services for Communications Service Providers, Worldwide”
    - **Leaders** (per Gartner, July 2015): Accenture, IBM
    - **SUPERFLUIDITY innovation**: provide more powerful platform that has the ability to homogeneously manage the network and its functions as a cloud platform across heterogeneous infrastructures all the way from the core to the access and vice-versa.

- “Cloud Infrastructure as a Service, Worldwide”
    - **Leaders** (per Gartner, May 2015): Amazon Web Services, Microsoft / Azure (Google, CenturyLink, VMware, IBM / SoftLayer are in the **Visionaries**)
    - **SUPERFLUIDITY innovation**: enable more powerful private and hybrid cloud offerings by empowering solutions to offer, transparently and with performance guarantees, the best of both bare-metal managed hosting and virtualised cloud environments spanning all the way towards the access of the network.

- “Public Cloud Infrastructure as a Service, Worldwide”
    - **Critical Capabilities** (per Gartner, October 2015)
        - Application Development: Amazon Web Services, Microsoft, Virtustream
        - Batch Computing: Amazon Web Services, Microsoft, Google, Rackspace
        - Cloud-Native Applications: Amazon Web Services, Microsoft, Google
        - General Business Applications: Amazon Web Services, Virtustream, Microsoft
    - **SUPERFLUIDITY innovation**: same as above

- “Public Cloud Storage Services”
    - **Leaders** (per Gartner, June 2015): Amazon Web Services, Microsoft, Google
    - **SUPERFLUIDITY innovation**: same as above

- “Cloud-Enabled Managed Hosting, Europe”
    - **Leaders** (per Gartner, June 2015): Rackspace, Interoute, BT Global Services, Claranet, Colt, CenturyLink, Verizon (Telefonica is in the **Niche Players**)
    - **SUPERFLUIDITY innovation**: same as above

- “Managed Hybrid Cloud Hosting, Europe”
    - **Leaders** (per Gartner, June 2016): Interoute, BT, Rackspace, Claranet, CenturyLink
    - **SUPERFLUIDITY innovation**: same as above

- “Data Center Networking”
    - **Leaders** (per Gartner, May 2016): Cisco, Arista (Dell and VMware are in the **Visionaries**, Brocade, Juniper and HPE in the **Challengers**, while NEC is in the **Niche Players**)
    - **SUPERFLUIDITY innovation**: will significantly advance the NFV and SDN state of the art, through introducing recursive network service (NS) modelling and decomposition, and more powerful/stateful SDN abstractions, based on extended finite state machines.

- “Network Equipment and Solutions Providers”
    - **Leaders** (per Gartner, August 2013): Alcatel-Lucent, Cisco, Ericsson, NEC, Nokia.
    - **SUPERFLUIDITY innovation**: offering a solution that spans over heterogeneous hardware architectures, offering performance- and scale-independent solutions in the cloud for networking deployments that are closely following the “on-demand” cloud computing innovations.

3 Exploitation Planning

Activities in this area aim to fulfil the Market Vision of the SUPERFLUIDITY project (please refer to section 3.3 of deliverable *D1.2 Project Vision and Roadmap*).

Section 3.1 outlines the exploitation plans of the individual partners. These were initialised from the Description of Work (DoW), but we revised them per the strategic or vision shifts of each partner.

Section 3.2 summarises the exploitation plan of the consortium as a whole. Other than to reflect consolidations (M&A), the overall plan has remained largely unchanged, which is a testimony to the insightfulness of project objectives, at the time of project definition.

As definition and implementation of 5G accelerates and the associated market segments continue to "pivot", we will keep the above exploitation plans aligned. Such changes will be reported, both internally, as reports of the Innovation Coordinator to the Management Board, as well as to the Project Officers, in the context of the regular reports and annual reviews.

3.1 Individual Partners

The exploitation plans of individual SUPERFLUIDITY partners are reported in Table 3 below.

PARTNER	EXPLOITATION PLANS
NOKIABLF	The SUPERFLUIDITY project brings significant contributions to the different Nokia solutions and products, from the application cloud management to the wireless part. The work conducted in SUPERFLUIDITY on Cloud RAN is highly valuable since it allows Nokia Bell-Labs to analyse and identify the right technological choice for the virtualization, hardware acceleration, service chaining, orchestration tool, in an advanced stage preparing the next iteration of the vRAN product. SUPERFLUIDITY helps Nokia Bell-Labs to realise and prototype the vision of a <u>cloud-based virtualised wireless access</u>. This exploitation will facilitate the <u>shift of additional applications from the cloud to the remote antenna and vice-versa</u>, <u>reducing the cost of the RRH and increasing the resource utilization at the cloud infrastructure</u>. Most importantly, such a solution will allow the <u>virtualization of the radio services</u> and thus the <u>convergence of all access services within a single framework</u>. In addition, the collaboration with partners of the SUPERFLUIDITY project, allows us to enhance our cloud RAN design by integrating the Mobile Edge Computing and evaluate the impact on Cloud RAN performance. This step is very important to prepare next product releases and evolution to meet the market needs. This work is initiated in the project in WP7 scope.

	Furthermore, the project results may improve multi-tenancy and multi-operator capability of future 5G Nokia products (e.g. Netact, Airframe, Next Cloud packet Core...) through the use of cloud, SDN and NFV technologies. SUPERFLUIDITY project will help Nokia to extend its technical and market position in the mobile networks market for the next 5G generation. The leading position can further be aided by securing the key concepts with IPRs to be made by the project and by influencing the standards forums.
NOKIAIL	NOKIAIL is targeting to exploit SUPERFLUIDITY's results to evolve NOKIA's portfolio with respect to NFV and MEC. First, extending NOKIA's NFV orchestration to better support C-RAN and MEC application is an immediate outcome of SUPERFLUIDITY. Then, SLA based placement across different POPs from the cloud edge to the core network is yet another enhancement to Cloudband VNF manager. It is expected, as envisioned by SUPERFLUIDITY, that future network functions would be composed of microservices that could be fastly instantiated, scaled or migrated. Providing the platform for those microservices, e.g., via containers and uni-kernels, is in the plans of Cloudband Infrastructure software. Additionally, extending the orchestration and management to support life cycle management in such a dynamic environment is in the roadmap of Cloudband Network director as well as Cloudband Application Manager.
BT	BT believes that in the next 2-3 years the worlds of cloud computing and virtualised network functions are coming together into a completely new model for telecommunications combining cloud services, network function virtualisation, and software defined networks. Entire enterprises will be able to run on virtualised servers, virtualised network functions, and virtualised network control services. The BT Vision is described by Howard Watson, our CEO of BT Technology, Service & Operations, at https://player.vimeo.com/video/166341115 (starts at 14 minutes) As part of this convergence, NFV has a critical role. We have already launched our first commercial service that uses NFV (Cloud Connect Intelligence), with more in the pipeline. The fast and flexible deployment and modification of services is what is most attractive about NFV. We are feeding the results from SUPERFLUIDITY into our internal OSS/BSS developments. Several technologies researched in SUPERFLUIDITY are of specific interest to BT. From an implementation point of view these include the work on extended finite state machines (since these hold the promise of having full defined and understood behaviour) and lightweight virtualisation techniques, such as containers. The approach /tool to verify composition also addresses an important issue. Once these techniques are more mature, we will consider whether they are suitable for our network.

	BT is also interested in the various demonstrations. For instance, currently we have an issue with CPE in the branches of, say, a bank – because site visits are expensive and hard to organise. Virtualisation offers the prospect of simplifying the equipment in the branch, by moving the more complex and maintenance-prone functions into a central site.
CITRIX	CITRIX is targeting to exploit SUPERFLUIDITY's results to evolve its Traffic Management appliances (the ByteMobile product family in general and the virtualised appliances belonging to such family in particular, e.g., the T3100 Adaptive Traffic Management Solution (see http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/bytemobile-adaptive-traffic-management-product-family.pdf, http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/bytemobile-virtual-adaptive-traffic-management-solution.pdf and the picture below). ByteMobile Adaptive Traffic Management Solution, Powered by ByteMobile Orchestration System 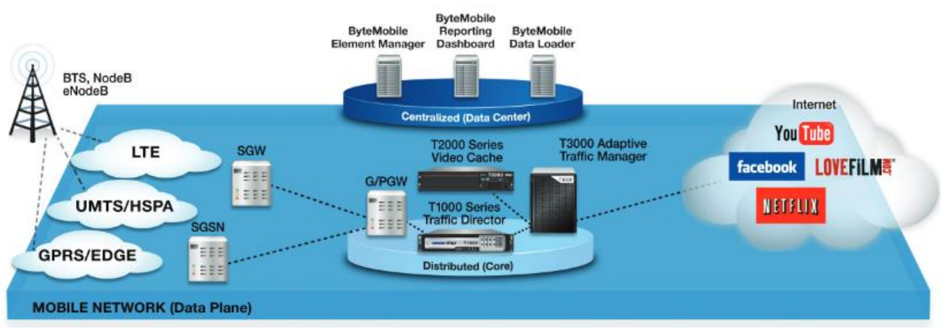 SUPERFLUIDITY's results will allow CITRIX to: • Extend the applicability of ByteMobile product family from the core network towards the wireless access by co-locating such (possibly virtualised) products next to wireless nodes (e.g., base stations) and taking advantage of radio network information to better optimise the traffic. • Improve the performance of the virtualised appliances (e.g., load balancing and application delivery) in order to offer a scalable service chaining orchestration of elementary functionalities with guaranteed SLA enforcement. Since a subset of the capabilities of the ByteMobile product family are being migrated to the NetScaler product family, the above exploitation plans may have to be adapted accordingly over the course of the project.
EBLINK	EBLINK's innovative Wireless Fronthaul solution represents a major technological advance in mobile network deployment, eliminating the last few hundred meters of fiber that are so costly for operators. EBLINK's solution reflects the evolution toward increasingly centralised telecom architectures, and allows for faster and easier deployment of 3G, 4G and future 5G networks. Wireless Fronthaul technology is considered crucial

	to the development of "Small Cells: pico and micro RHHs", to ensure needed coverage and high-speed service. EBLINK will exploit SUPERFLUIDITY's results in order to trial his solutions together with vendors such as Nokia, and mobile networks operators as British Telecom and Telefonica to advance their products and make them better fitting to their future needs. EBLINK plans to embed the SUPERFLUIDITY results in the management of heterogeneous access networks in their future product plan with the target to help mobile network operators to densify their network by implementing backhaul/fronthaul with low latency that will be compatible with access traffic increase.
INTEL	As NFV and SDN are expected to be foundational technologies in the realisation of 5G networks that are elastic, programmable and dynamically manageable it is clear that certain areas such as enhanced platform awareness (EPA), workload affinity for platform features, intelligence orchestration and approaches to resource description models will play a critical role. Research in WPs 4-6 into these topics within SUPERFLUIDITY is mapping directly to Intel's internal research agenda. To ensure convergence between technologies being developed by Intel Business Units (BUs) and external customer needs, we have aligned the efforts of our activities within SUPERFLUIDITY with our internal research framework Apex Lake which is focused on realising intelligent orchestration proof of concepts. Intel is actively engaged in both open standards and commercial software enablement. Intel has a strong involvement in a number of key open source projects aligned to NFV and SDN including OpenStack, OpenDaylight, Open vSwitch, OpenFlow and the Data Plane Development Kit. In addition, Intel is a leading member of the Mobile Edge Computing industry initiative, which is focused on opening up the radio access interface for services and applications in the radio access network (RAN). Intel will use SUPERFLUIDITY as proof point for open source technologies, MEC and the value of an open source ecosystem to ensure interoperability for 5G deployments and associated use cases. From a commercial perspective, Intel's exploitation plans in SUPERFLUIDITY lie in the realisation that convergence of fixed and mobile technologies presents the opportunity for sharing of infrastructure resources for efficient service delivery. Intel is therefore looking forward to bring SUPERFLUIDITY results to the field so as to enable an efficient and highly manageable infrastructure landscape, making use of Standard High Volume (SHV) elements to empower XaaS (Platform or Infrastructure as a Service) solutions with resource pooling, on-demand scaling of capacity and allocation of resources proportionate to revenue.

<table>
<tr>
<td>NEC</td>
<td>
NEC is targeting to exploit SUPERFLUIDITY's results to evolve its solution portfolio in three main areas:

Small cell solution (both outdoor and in-building) (http://www.nec.com/en/global/solutions/nsp/sc/index.html)
LTE-Advanced high-capacity base station equipment (http://www.nec.com/en/press/201302/global_20130227_04.html)
Cloud solution (http://www.nec.com/en/global/solutions/cloud/index.html?)

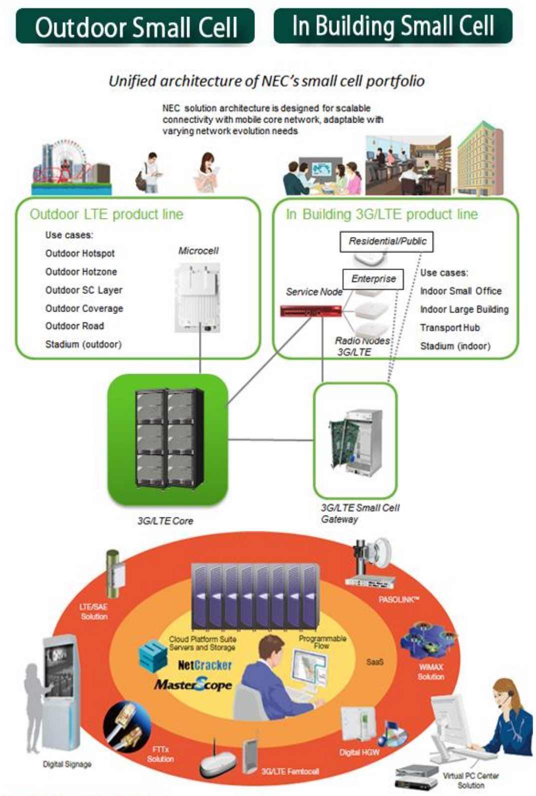

Figure 5 NEC's Solutions for Carrier Cloud Services
In <u>outdoor scenarios</u>, NEC will target the technology transfer to its business units responsible for the "New Last Mile" concept: the one-stop-shop to help mobile network operators deploy the next layer of LTE capacity and coverage in a rapid, economical and integrated way and keep pace with skyrocketing data demand. Such solution is designed to simplify the end-to-end process of delivering the LTE brand promise in complex 3D urban landscapes, VIP customer sites outside the current macro footprint and rural areas in line with their spectrum licensing obligations. The initiative leverages NEC's leadership in open, multivendor HetNet technology and the most comprehensive small cell and backhaul portfolio in the market.
In <u>in-building scenarios</u>, we target a technology transfer to NEC's Enterprise RAN (E-RAN) solution that already utilise innovative multimode LTE and 3G coverage for enterprise internal phone network.
</td>
</tr>
</table>

	As for LTE-Advanced High-Capacity Base Station Equipment, NEC is planning to evolve its base station from a system where specialist processing is done to a modular, open solution integrated in an ecosystem of changeable components as the "evolution of the base station concept". Overall, NEC's exploitation plan for SUPERFLUIDITY is to integrate the three above areas in a horizontal manner and provide a true end-to-end solution with mobile edge to cloud cooperation (seamless and quality-aware migration of service and applications from radio head to cloud and vice-versa).
ONAPP	OnApp intend to exploit SUPERFLUIDITY's results to its product portfolio in three key areas: • Increased storage performance: SUPERFLUIDITY results will be used for the OnApp Storage (http://onapp.com/platform/onapp-storage/); • Rapid and scalable service provisioning: using low overhead resource management systems generated in SUPERFLUIDITY, OnApp intend to improve its CDN (http://onapp.com/platform/onapp-cdn/) and Federated Market products (http://onapp.com/federation/). Thin Hypervisor platforms can more easily be powered up and migrated; • Energy usage reduction for modern hardware architectures: extending network architectures to the edge to make intelligent use of micro-servers (e.g., continuing the work of OnApp and TUD in the Euroserver FP7 project - http://www.euroserver-project.eu/) that will be part of the next generation of OnApp hardware platforms. • Improvements to the OnApp User Interface / dashboard to cope with and manage the expected scale of virtualised functions. The main motivation for OnApp to exploit SUPERFLUIDITY results is the belief that moving services towards the end-user is a key trend that is visible from the Cloud and Datacenter perspective. OnApp has a huge client base (3,000+ clouds, 900+ service providers in 93 countries) that use the OnApp CDN platform to expand their infrastructure beyond their own physical hardware, to bring content closer to their end-users. In addition, the OnApp Federation (that was a result of the FP7 Trilogy2 project where OnApp worked together with Intel, NEC and BT, http://trilogy2.eu/) is being used to allow customers to provision computation near to end-users. Adoption of the Federation and CDN platform highlight the requirement from customers to bring services close to customers. The current system requires a-priori provisioning of full Virtual Machines. By reducing the size of Virtual Machines as being proposed in SUPERFLUIDITY, it will allow for more rapid, dynamic provisioning. The optimisation of the Hypervisor to allow higher performance between guests will enable a more fluid use of resources available throughout the infrastructure. This combined with lighter VMs will allow better

	dynamicity of services to respond, rapidly to end-user demands and reduce the amount of over-provisioning of resources currently needed to be able to respond to demand.
ALB	The ALB exploitation plan intends to transfer the outputs of the SUPERFLUIDITY project into the business units (BU), enhancing the ALB's products and services portfolio in the long-run. In particular, ALB is deeply involved in the Mobile Edge Computing (MEC) technology, following and contributing to the ongoing standardization work. Within the scope of the SUPERFLUIDITY project, ALB plans to develop and implement a MEC prototype in order to validate the concept and show the edge computing benefits. This Proof-of-Concept (PoC) will allow ALB to acquire know-how and eventually establish the foundations for a future ALB MEC solution. As the MEC technology is still in early stages of development, ALB believes that can take a significant advantage regarding other competitors, as today full edge computing solutions are not available in the market. The relationship of MEC with other important technologies like SDN and NFV (e.g. C-RAN, EPC) is also an area where ALB intends to acquire know-how, validating and prototyping this integration. The MEC technology will help 5G to achieve one of the key parameter indicators: the end-to-end-latency < 1ms. Considering that the laws of physics cannot be overtaken, the only way to continue reducing latency to magnitudes of 1ms (end-to-end) is getting the services closer to the end users. Recently, ETSI has created the MEC (Mobile Edge Computing) ISG in order to cover this issue. At this moment, this group has produced a basic architecture and some guidelines about how edge computing should work. SUPERFLUIDITY, via ALB, is actively contributing to this group. Although the MEC concept is not tied to any particular technology, it benefits with the introduction of C-RAN (Centralised/Cloud-RAN) in the network. Currently, the deployment of a cloud-based Application on every site/antenna (eNB) is not economically feasible, as even medium sized countries may have tens of thousands of spots. However, with the C-RAN advent, this value may get reduced to a few tens, making it more suitable. The SUPERFLUIDITY project will integrate C-RAN and MEC, showing that together, ultra-low latency values and bandwidth efficiencies, can be achieve at a reasonable cost. Furthermore, both technologies can share the cloud resources, taking advantage of management synergies and reducing overall costs. The SUPERFLUIDITY project will implement the MEC architecture from scratch. It intends to fully be able to manage and orchestrate the whole MEC solution by the end of the project. In addition to the ETSI MEC

	architecture, the SUPERFLUIDITY implementation will have the following key properties: • Compatibility with existing mobile networks - there is no need for network changes; the network is MEC-unaware. • Compatibility with existing applications - applications don't need to change to provide services, unless they want to take advantage of MEC Services (APIs). • Transparency to end users - the user is not aware about whether the service is provided from the core or the edge (they just feel different latencies). • Integration with NFV infrastructure - the operator can take advantage of NFV and MEC synergies, reducing overall costs.
REDHAT	REDHAT is planning to use SUPERFLUIDITY results to improve its "Red Hat Enterprise Linux OpenStack Platform" (see https://access.redhat.com/products/red-hat-enterprise-linux-openstack-platform), in particular the OVA-R (Open Virtual Appliance for Red Hat Enterprise Linux OpenStack Platform, see https://access.redhat.com/products/red-hat-enterprise-linux-openstack-platform/ova-r). These solutions have their roots in the Open Virtualization Alliance, OVA, see https://openvirtualizationalliance.org/ and are linked to the Red Hat Enterprise Virtualization initiative (http://www.redhat.com/promo/linux-virtualization/). Such platform delivers an integrated foundation to create, deploy, and scale a secure and reliable public or private OpenStack cloud. Red Hat Enterprise Linux OpenStack Platform combines the world's leading enterprise Linux and the fastest-growing cloud infrastructure platform to give customers the ability to scale and quickly meet customer demands without compromising on availability, security, or performance. SUPERFLUIDITY results will enhance the capabilities of such platform, currently focused on cloud developments, with the possibility to use and deploy virtualised functions across different appliances located in the wireless access domain next to e.g., base stations.
TELCA	TELCARIA is a young SME dealing with network service virtualization for network operators. SUPERFLUIDITY technologic achievements will allow exploit the new ecosystem created by SUPERFLUIDITY by developing and bringing to the market innovative and ground-breaking services and applications that can take advantage of the information on radio network capabilities and conditions available at the base station. This will enable TELCARIA to do rapid prototyping, design validation and direct deployment on its customers at a much faster pace than traditional vendors thus enabling gaining of market share and creation of new jobs. One of the main areas in which TELCARIA foresees exploitation potential of SUPERFLUIDITY is the advances in technologies for integrating

	heterogeneous wireless networks (up to RAN level) and in architectures to optimise the reuse of functionality across heterogeneous access technologies for 5G. In this context, Telcaria's team has participated in numerous collaborations with Indra, Telefonica, Huawei, NEC, the Spanish Ministry of Defence, etc., and we believe SUPERFLUIDITY will help strengthen future collaborations and contracts as new services in relation with 5G.
TID	Telefonica recently (at MWC 2014, see http://www.telecompaper.com/news/telefonica-unveils-unica-virtualization-infrastructure--998081) unveiled a network infrastructure initiative called UNICA. UNICA is a pioneering infrastructure that allows to make use of virtualised functions in the core of the network. Telefonica's target is to have more than 30 per cent of new infrastructure virtualised by 2016. The initiative started pilot operations in June 2014 and has the target to provide Telefonica with a framework for its global, end-to-end virtualization deployment. Telefonica has been running trials with several vendors (including Nokia and NEC) covering test cases related to e.g., cloud deployment of services and compliance of OpenStack infrastructure. TID will transfer the SUPERFLUIDITY results to the business units in charge of UNICA, and perform trial demonstrations of them, in order to achieve a direct availability of services based on the project results on the UNICA infrastructure. Furthermore, Telefonica is currently working in the application of NFV technologies to the home environment and is running a home router virtualization (vCPE) trial in Brazil prior to launching the commercial service, planned for the first quarter of 2015, as a first step towards a progressive deployment in all the operators of the Telefonica Group. The vCPE concept allows moving all software and control functions at the home environment into the cloud, avoiding the need of installing complex elements at homes and facilitating the seamless update of existing services and the deployment of new ones. TID, as current coordinator of the trial, will be in the position of including wireless access aspects to the vCPE services by applying the SUPERFLUIDITY results.
USTR	SUPERFLUIDITY will provide USTR with a unique opportunity to improve the streaming/muxing software within CDN setups at the edge of the network. The targeted product will be in particular the UNIFIED REMIX, USTR product that creates personalised streams based on a recommendation, a ruleset or even an external repository such as an ad network. The opportunity of using SUPERFLUIDITY's results is to mux content as late as possible has with the benefit of reducing latency and lower the amount of CDN internal traffic thus reducing cost at operations level. USTR will bring its A/V streaming expertise to the project as well as access to the professional streaming solutions it creates and provide use-case testing and integration. More specific exploitation plans include:

	Location awareness created by virtualised Mobile Edge Computing support in SUPERFLUIDITY: • improvement of streaming/muxing software in the edge (scripts to deploy unified streaming software in the cloud) • remix, personalization, enabling costumers better opportunities for monetization • transmuxing in realistic virtualised edge computing, reducing backend traffic and caching requirements • better reception/video optimization based on Radio Access information • these will all reduce network transmission and caching cost for customers running our software The scale independence offered by the Nokia CloudBand and the intel telemetry setup: • enable a video streaming origin that runs in a private cloud and scales compute/network resources on the go automatically • This will benefit Unified Streaming customers that currently typically runs software on public clouds or private servers NEC work on fast instantiation and minimalistic kernels: • will enable unified streaming to run on as minimalistic kernels as possible, enabling USTR to offer highly cost efficient implementations of unified origin in a cloud environment.
ACADEMIC PARTNERS	Academic partners plan the exploitation of SUPERFLUIDITY results in order to further increase their standings and ranking in the academic area through participation to key scientific conferences. Furthermore, research results from SUPERFLUIDITY will flow into the academic curricula by endorsing solutions and new paradigms in teaching programmes. Both MSc. and PhD students will be targeted. Teaching material will be updated with the latest technological achievements so as to improve the quality of the offered studies. Last but not least, academic partners will use SUPERFLUIDITY's results for attracting industry consulting contracts and to launch new innovative software-based academic spin-offs. As a matter of fact, members of CNIT involved in SUPERFLUIDITY just founded a spin-off, together with some talented, newly graduated, engineers. The spin-off is already selling mobile applications developed for the health sector and it is developing applications for the cultural heritage and tourism industry. The participation to SUPERFLUIDITY could offer significant opportunities to the spin-off, which focuses on the development of applications using cutting edge mobile technologies. For example, the capability of doing efficient real time streaming and sharing of multimedia information on mobile devices will be factored in future products of the spin-off.

Table 3: Per-partner Exploitation Plans

3.2 Consortium as a Whole

The overall exploitation plan of the consortium as a whole is reported in Table 4 below.

Exploitation Plan: Consortium as a Whole

In terms of consortium, SUPERFLUIDITY is bringing partners with complementary expertise addressing different roles in the exploitation of project results, here we present the overall vision of the consortium starting from bottom up (from hardware, through operating system up to services and their orchestration including the architectural aspects):
• INTEL: standard high volume elements (client, edge and server platforms, networking equipment)
• NEC: integration of software solutions for small cells, LTE base stations and cloud resources
• NOKIA
    o NOKIABLF: enhancements of the remote radio head (RRH); preparation of the next version of the vRAN
    o NOKIAIL: management tools for orchestrating virtualised functions across heterogeneous access
• REDHAT: operating systems and orchestration running on top of standard high volume (SHV) elements
• EBLINK: wireless fronthaul solution for improved coverage of heterogeneous access
• TELCARIA: integrate heterogeneous wireless networks information for running innovative services
• CITRIX: virtualised traffic management spreading over the converged infrastructure (up to the access)
• ONAPP: virtualised storage services spreading over the converged infrastructure (up to the access);
• UNIFIED STREAMING: streaming services spreading over the converged infrastructure (up to the access);
• TELECOM OPERATORS (TID, ALTICE LABS and BRITISH TELECOM): converged infrastructure moving service functionalities between core and access.

Table 4: Exploitation Plan of the Consortium as a Whole

4 Conclusion

As reported by this deliverable, the SUPERFLUIDITY project is exhibiting progress in terms of planning, and even initiating execution, of activities in the areas of innovation and commercial exploitation.

Progress in executing the project's Innovation and Exploitation Plan will be subsequently reported in:

- D8.5 First Report on Innovation and Exploitation Actions (M20)
- D8.8 Final Report on Innovation and Exploitation Actions (M30)

5 References

[8021Q] IEEE Std. 802.1Qbg-2012, "IEEE Standard for Local and Metropolitan Area Networks — Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks— Amendment 21: Edge Virtual Bridging," July 5, 2012, http://standards.ieee.org/getieee802/download/802.1Qbg-2012.pdf, p. 191.

[ARN2010] Applied Technology and Innovation Management, Heinrich Arnold, Michael Erner, Peter Möckel, Christopher Schläffer Editors, Springer Heidelberg Dordrecht London New York

[ASG2016] H. J. Asghar, L. Melis, C. Soldani, E. De Cristofaro, M. A. Kaafar, and L. Mathy: "SplitBox: Toward Efficient Private Network Function Virtualization" ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization (HotMiddlebox 2016), Co-located with ACM SIGCOMM'16, August 22-26, 2016, Salvador Florianópolis, Brazil.

[ATT2016] AT&T "ECOMP (Enhanced Control, Orchestration, Management & Policy) Achitecture" White Paper, March 2016

[BAN12] M. Bansal, J. Mehlman, S. Katti, and P. Levis, "OpenRadio: A Programmable Wireless Dataplane," in Proc. of HotSDN'12, Aug. 2012.

[BAN15] Manu Bansal, Aaron Schulman, and Sachin Katti. 2015. Atomix: a framework for deploying signal processing applications on wireless infrastructure. In Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation (NSDI'15). USENIX Association, Berkeley, CA, USA, 173-188.

[BEL2010] Beloglazov, Anton, and Rajkumar Buyya. "Energy efficient resource management in virtualized cloud data centers." 10th IEEE/ACM Int. Conf. on Cluster, Cloud and Grid Computing. IEEE Computer Society, 2010.

[BHA2013] Bhatnagar, Vaibhav, et al. "An FPGA Software Defined Radio Platform with a High-Level Synthesis Design Flow." Vehicular Technology Conference (VTC Spring), 2013 IEEE 77th. IEEE, 2013.

[BHAN10] G. Bhanage, I. Seskar, R. Mahindra, and D. Raychaudhuri, "Virtual Basestation: Architecture for an Open Shared WiMAX Framework," in Proc. of VISA'10, 2010.

[BIA12] D. Garlisi, F. Giuliano, I. Tinnirello, P. Gallo, F. Gringoli and G. Bianchi, "Deploying virtual MAC protocols over a shared access infrastructure using MAClets," Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on, Turin, 2013, pp. 17-18.

[BIA14] Giuseppe Bianchi, Marco Bonola, Antonio Capone, Carmelo Cascone: OpenState: programming platform-independent stateful openflow applications inside the switch. ACM Comp. Commun. Rev. 44(2): 44-51 (2014)

[BiaArx14] Bianchi, G., Bonola, M., Picierro, G., Pontarelli, S., & Monaci, M. (2013). StreaMon: a data-plane programming abstraction for Software-defined Stream Monitoring. arXiv preprint arXiv:1311.2442.

[BOS13] P. Bosshart, G. Gibb, H.-S. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz, “Forwarding metamorphosis: Fast programmable match-action processing in hardware for sdn,” in ACM SIGCOMM 2013. ACM, 2013, pp. 99–110.

[BOT2012] Botero, Juan Felipe, et al. "Energy efficient virtual network embedding." Communications Letters, IEEE 16.5 (2012): 756-759.

[CHE15] A. Checko et al., "Cloud RAN for Mobile Networks—A Technology Overview," in IEEE Communications Surveys & Tutorials, vol. 17, no. 1, pp. 405-426, Firstquarter 2015.

[COH2016] Alejandro Cohen, Asaf Cohen and Omer Gurewitz, "Secure Group Testing", IEEE International Symposium on Information Theory (ISIT), Barcelona, July 2016.

[Cra13] B. Mack-Crane. “OpenFlow Extensions”. In: US Ignite ONF GENI workshop, October 8, 2013.

[DET2013] Gregory Detal, Benjamin Hesmans, Olivier Bonaventure, Yves Vanaubel, and Benoit Donnet. 2013. Revealing middlebox interference with tracebox. In Proceedings of the 2013 conference on Internet measurement conference (IMC '13).

[DET2016] Andrea Detti, Bruno Ricci, Nicola Blefari-Melazzi: “Tracker-assisted rate adaptation for MPEG DASH live streaming”. To appear in the Proceedings of the IEEE International Conference on Computer Communications (INFOCOM) 2016. 2016, April 10-15. San Francisco, United States of America.

[DOB2009] M. Dobrescu, N. Egi, K. Argyraki, B.-G. Chun, K. Fall, G. Iannaccone, A. Knies, M. Manesh and S. Ratnasamy, ”RouteBricks: Exploiting Parallelism to Scale Software Routers”, ACM SOSP'09, 2009.

[DON2012] Dong, Yaozu, et al. "High performance network virtualization with SR-IOV." Journal of Parallel and Distributed Computing 72.11 (2012): 1471-1480.

[DPDK2016] "DPDK: Data Plane Development Kit – What it is". dpdk.org. Retrieved June 2016.

[DUN09] R. Duncan and P. Jungck, “PacketC: Language for High-Performance Packet Processing”, IEEE HPCC'09, 2009.

[ETT2016] G. Bianchi, E. Biton, N. Blefari-Melazzi, I. Borges, L Chiaraviglio, Pedro de la Cruz Ramos, P. Eardley, F. Fontes, M. J. McGrath, L. Natarianni, D. Niculescu, C. Parada, M. Popovici, V. Riccobene, S. Salsano, B. Sayadi, J. Thomson, C. Tselios, G. Tsolis, ”SUPERFLUIDITY: A Flexible Functional Architecture for 5G Networks”, Article accepted for publication on Transactions on Emerging Telecommunications Technologies, Wiley.

[FAWG] Openflow Forwarding Abstractions Working Group Charter. http: //goo.gl/TtLtw0. Apr. 2013.

[FOS11] Foster, Nate, et al. "Frenetic: A network programming language." ACM SIGPLAN Notices. Vol. 46. No. 9. ACM, 2011.

[GUD2013] Aditya Gudipati, Daniel Perry, Li Erran Li, Sachin Katti, “SoftRAN: Software Defined Radio Access Network.” ACM HotSDN’13, August 16, 2013, Hong Kong, China.

[HAN2010] S. Han, K. Jang, K. Park, and S. Moon, “PacketShader: a GPU-Accelerated Software Router”, ACM SIGCOMM'10, 2010.

[HON2011] Michio Honda, Yoshifumi Nishida, Costin Raiciu, Adam Greenhalgh, Mark Handley, and Hideyuki Tokuda. 2011. Is it still possible to extend TCP?. In Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference (IMC '11).

[HUA14] C. L. I, J. Huang, R. Duan, C. Cui, J. (. Jiang and L. Li, "Recent Progress on C-RAN Centralization and Cloudification," in IEEE Access, vol. 2, no. , pp. 1030-1039, 2014.

[HWA2015] J. Hwang, K. K. Ramakrishnan and T. Wood, "NetVM: High Performance and Flexible Networking Using Virtualization on Commodity Platforms," in IEEE Transactions on Network and Service Management, vol. 12, no. 1, pp. 34-47, March 2015.

[JEY13] V. Jeyakumar, M. Alizadeh, C. Kim, and D. Mazieres, “Tiny packet programs for low-latency network control and monitoring,” in ACM Workshop on Hot Topics in Networks (HOTNETS 2013), 2013.

[JEY2013] Vimalkumar Jeyakumar, Mohammad Alizadeh, David Mazières, Balaji Prabhakar, Changhoon Kim, and Albert Greenberg. 2013. EyeQ: practical network performance isolation at the edge. In Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation

[KAM2016] Joseph Kampeas, Asaf Cohen and Omer Gurewitz, "On Secrecy Rates and Outage in Multi-User Multi-Eavesdroppers MISO Systems",IEEE International Symposium on Information Theory (ISIT), Barcelona, July 2016.

[KOH2000] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, “The Click Modular Router”, ACM Transactions on Computer Systems 18(3), Aug 2000.

[KOK10] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, “NVS: a virtualization substrate for WiMAX networks,” in Proc. of MobiCom’10, 2010.

[LAC2013] Katrina LaCurts, Shuo Deng, Ameesh Goyal, and Hari Balakrishnan. 2013. Choreo: network-aware task placement for cloud applications. In Proceedings of the 2013 conference on Internet measurement conference (IMC '13).

[MAD2013] Unikernels: Library operating systems for the cloud. A. Madhavapeddy, R. Mortier, C. Rotsos, D. Scott, B. Singh, T. Gazagnaire, S. Smith, S. Hand, and J. Crowcroft. SIGPLAN Not., 48(4):461–472, Mar. 2013.

[MAE11] Maeder, A., Mancuso, V., Weizman, Y., Biton, E., Rost, P., Perez-Costa, X., & Gurewitz, O., “FLAVIA: Towards a generic MAC for 4G mobile cellular networks”, Future Network & Mobile Summit, 2011

[MAI2012] S. Maital, D. V. R. Seshadri, Innovation Management, Response Books

[MAR2014] ClickOS and the Art of Network Function Virtualization. J. Martins, M. Ahmed, C. Raiciu, V. Olteanu, M. Honda, R. Bifulco, F. Huici. NSDI 2014.

[MEC] “Mobile Edge Computing – Introductory Technical White Paper”. https://portal.etsi.org/Portals/0/TBpages/MEC/Docs/Mobile-edge_Computing_-_Introductory_Technical_White_Paper_V1%2018-09-14.pdf

[MEK2016a] R. Mekuria, J. Fennema, R. Belleman, D. Griffioen, USTR – Unified Streaming: “Multi-Protocol Video Delivery with Late Trans-Muxing” accepted paper at ACM Multimedia, 15 - 19 October 2016, Amsterdam, The Netherlands.

[MEK2016b] Rufael Mekuria, Michael McGrath, Christos Tselios, Dirk Griffioen, George Tsolis, Shahar Beiser: “KPI Mapping for Virtual Infrastructure Scaling for a Realistic Video Streaming Service Deployment”, 8th International Conference on Quality of Multimedia Experience (IEEE QoMEX), Lisbon, Portugal, June 6-8, 2016.

[Mey13] D. Meyer, “Openflow: Today’s reality, tomorrow’s promise? an architectural perspective,” in SDN Summit 2013, Paris, France, March 2013, 2013.

[MIJ2016a] R. Mijumbi, J. Serrat, J. L. Gorricho, N. Bouten, F. De Turck and R. Boutaba, "Network Function Virtualization: State-of-the-Art and Research Challenges," in IEEE Communications Surveys & Tutorials, vol. 18, no. 1, pp. 236-262, Firstquarter 2016.

[MIJ2016b] R. Mijumbi, J. Serrat, J. l. Gorricho, S. Latre, M. Charalambides and D. Lopez, "Management and orchestration challenges in network functions virtualization," in IEEE Communications Magazine, vol. 54, no. 1, pp. 98-105, January 2016.

[MON13] Monsanto, Christopher, et al. "Composing Software Defined Networks." NSDI. 2013.

[MOR99] Robert Morris, Eddie Kohler, John Jannotti, and M. Frans Kaashoek. 1999. The Click modular router. SIGOPS Oper. Syst. Rev. 33, 5 (December 1999), 217-231. DOI=10.1145/319344.319166 http://doi.acm.org/10.1145/319344.319166

[MOS14] M. Moshref, A. Bhargava, A. Gupta, M. Yu, and R. Govindan. Flow-level State Transition as a New Switch Primitive for SDN. In Proceedings of the ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN’14), August 2014.

[MUE2010] Mueck, Markus, et al. "ETSI reconfigurable radio systems: status and future directions on software defined radio and cognitive radio standards." Communications Magazine, IEEE 48.9 (2010): 78-86.

[MWH2014] I. Marinos, R. N. M. Watson, and M. Handley. Network Stack Specialization for Performance. In Proceedings of the ACM SIGCOMM 2014 Conference, pages 175–186, 2014.

[NFV001] ETSI GS NFV-SWA 001 V1.1.1 “Network Functions Virtualisation (NFV); Virtual Network Functions Architecture”, December 2014

[NFV007] ETSI GS, “Network Functions Virtualisation (NFV); Infrastructure; Methodology to describe Interfaces and Abstractions”; NFV-INF 007 V1.1.1 (2014-10)

[NFV2014] ETSI Network Functions Virtualisation, http://www.etsi.org/technologies-clusters/technologies/nfv

[NGMN15] NGMN Alliance “Further Study on Critical C-RAN Technologies”, White Paper, March 31, 2015

[NSM2016] “Hybrid IP/SDN networking: open implementation and experiment management tools”, S. Salsano, P. L. Ventre, F. Lombardo, G. Siracusano, M. Gerola, E. Salvadori, M. Santuari, M. Campanella, L. Prete. Article accepted for publication on IEEE Transactions on Network and Service Management.

[NVGRE] M. Sridharan et al., “NVGRE: Network Virtualization Using Generic Routing Encapsulation,” IETF Draft- sridharan-virtualization-nvgre-03.txt, Aug. 2013

[NVO] Narten et al., “Problem Statement: Overlays for Network Virtualization,” IETF Draft draft-ietf-nvo3-overlay-problem-statement-04, July 31, 2013

[ODPL] OpenDataPlane. http://www.opendataplane.org

[OEC] OpenEdgeComputing, available online: http://openedgecomputing.org/index.html

[OF1.4] Open Networking Foundation “Openflow specification, ver 1.4”, October 2013

[OSF2016] OpenStack Foundation “Accelerating NFV Delivery with Openstack”, Technical Report, February 2016

[PAL2015] Shoumik Palkar, Chang Lan, Sangjin Han, Keon Jang, Aurojit Panda, Sylvia Ratnasamy, Luigi Rizzo, and Scott Shenker. 2015. E2: a framework for NFV applications. In Proceedings of the 25th Symposium on Operating Systems Principles (SOSP '15). ACM, New York, NY, USA, 121-136. DOI=http://dx.doi.org/10.1145/2815400.2815423

[PEN13] K. Pentikousis, Y. Wang, and W. Hu, “Mobileflow: Toward software defined mobile networks,” Communications Magazine, IEEE, vol. 51, no. 7, pp. 44–53, 2013.

[Per13] P. Peresini, M. Kuzniar, and D. Kostic. “OpenFlow Needs You! A Call for a Discussion About a Cleaner OpenFlowAPI”. In Proc. of the EU Workshop on Software Defined Network (EWSDN). 2013.

[PICSIG] PCI-SIG, “Single Root I/O Virtualization and Sharing 1.1 Specification,” http://www.pcisig.com/members/down-loads/specifications/iov/sr-iov1_1_20Jan10.pdf, available only to members.

[POP2010] Lucian Popa, Minlan Yu, Steven Y. Ko, Sylvia Ratnasamy, and Ion Stoica. 2010. CloudPolice: taking access control out of the network. In Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks (Hotnets-IX)

[POP2012] Lucian Popa, Gautam Kumar, Mosharaf Chowdhury, Arvind Krishnamurthy, Sylvia Ratnasamy, and Ion Stoica. 2012. FairCloud: sharing the network in cloud computing. In Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication (SIGCOMM '12)

[POP2013] Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeffrey C. Mogul, Yoshio Turner, and Jose Renato Santos. 2013. ElasticSwitch: practical work-conserving bandwidth guarantees for cloud computing. SIGCOMM Comput. Commun. Rev. 43

[QAZ2013] Zafar Ayyub Qazi, Cheng-Chun Tu, Luis Chiang, Rui Miao, Vyas Sekar, and Minlan Yu. 2013. SIMPLE-fying middlebox policy enforcement using SDN. SIGCOMM Comput. Commun. Rev. 43, 4 (August 2013)

[RAD15] Radisys Corp. “Evaluating Cloud RAN Implementation Scenarios” White Paper, June 2015

[RAI2012] Costin Raiciu, Mihail Ionescu, and Dragoş Niculescu. 2012. Opening up black box networks with CloudTalk. In Proceedings of the 4th USENIX conference on Hot Topics in Cloud Computing (HotCloud'12).

[RAM2009] Ram, Kaushik Kumar, et al. "Achieving 10 Gb/s using safe and transparent network interface virtualization." Proc. 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments. ACM, 2009.

[RFC6325] R. Perlman et al., "Routing Bridges (RBridges): Base Protocol Specification," IEEE RFC 6325, July 2011, 99 pages, http://tools.ietf.org/html/

[RIZ2012] L. Rizzo, "Netmap: a novel framework for fast packet i/o. In Proceedings of the 2012 USENIX conference on Annual Technical Conference", USENIX ATC'12, USENIX Association, Berkeley, CA, USA, 2012.

[RIZ2012a] L. Rizzo, M. Carbone, G. Catalli, "Transparent acceleration of software packet forwarding using netmap", in Proceedongs of IEEE INFOCOM'12, 2012.

[SACH08] J. Sachs and S. Baucke, "Virtual radio: a framework for configurable radio networks," in Proc. of WICON'08, Nov. 2008.

[SAT15] Satyanarayanan, M., Schuster, R., Ebling, M., Fettweis, G., Flinck, J., Joshi, K., Sabnani, K. "An Open Ecosystem for Mobile-Cloud Convergence", in IEEE Communications Magazine, Volume 53, Number 3, March 2015

[SEK2012] Design and implementation of a consolidated middlebox architecture. V. Sekar, N. Egi, S. Ratnasamy, M. K. Reiter, and G. Shi. 9th USENIX conf. on Networked Systems Design and Implementation, NSDI 2012.

[SHE2012b] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, V. Sekar, "Making middleboxes someone else's problem: network processing as a cloud service", SIGCOMM , 2012.

[SIR2016] Giuseppe Siracusano, Roberto Bifulco, Simon Kuenzer, Stefano Salsano, Nicola Blefari Melazzi, Felipe Huici: "On-the-Fly TCP Acceleration with Miniproxy", ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization (HotMiddlebox 2016), Co-located with ACM SIGCOMM'16, August 22-26, 2016, Salvador Florianópolis, Brazil.

[SON13] H. Song, "Protocol-oblivious forwarding: Unleash the power of sdn through a future-proof forwarding plane," 2nd ACM SIGCOMM HotSDN Workshop, 2013, pp. 127–132.

[STO2016a] Radu Stoenescu, Matei Popovici, Lorina Negreanu, and Costin Raiciu, UPB – University Politehnica of Bucharest: "Scalable Symbolic Execution for Modern Networks", ACM SIGCOMM 2016, Florianópolis, Brazil, August 22-26, 2016.

[STO2016b] Radu Stoenescu, Dragos Dumitrescu, Costin Raiciu: "OpenStack Networking for Humans: Symbolic Execution to the Rescue" invited paper, IEEE LANMAN 2016, Rome, Italy, June 13 -15, 2016.

[STT] B. Davie, Ed., J. Gross, "A Stateless Transport Tunneling Protocol for Network Virtualization (STT)," IETF Draft draft-davie-stt-03.txt, Mar. 12, 2013, 19 pages, http://tools.ietf.org/html/draft-davie-stt-03

[TAN11] K. Tan et al., “Sora: High-Performance Software Radio Using General-Purpose Multi-Core Processors,” Communications of the ACM, vol. 54, pp. 99–107, Jan. 2011.

[TIN12] I. Tinnirello, G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, “Wireless MAC Processors: Programming MAC Protocols on Commodity Hardware”, IEEE INFOCOM, March 2012.

[TSE2016] C. Tselios and G. Tsolis "On QoE-awareness Using Virtualized Probes in 5G Networks", submitted for publication in IEEE CAMAD 2016, October 20-24, Toronto, Canada.

[VXLAN] M. Mahalingam et al. “VXLAN: A Framework for Over-laying Virtualized Layer 2 Networks over Layer 3 Net-works,” IETF Draft draft-mahalingam-dutt-dcops-vxlan-04.txt, May 8, 2013.

[VZN2016] Verizon “SDN-NFV Reference Architecture v1.0” Verizon Network Infrastructure Planning, February 2016

[WU2013] NON Existent Paper – needs to be deleted from the text as well

[WU2014] Wu, Wenfei, et al. "PRAN: Programmable Radio Access Networks." 13th ACM HotNet Workshop, 2014.

[Yad11] N. Yadav and D. Cohn. “OpenFlow Primitive Set”, available online at http://goo.gl/6qwbg. July 2011.

[YAP10] K.-K. Yap et al., “Blueprint for introducing innovation into wireless mobile networks,” Proc. of VISA’10, 2010.

[ZAK10] Y. Zaki, L. Zhao, C. Grg, A. Timm-Giel, “A Novel LTE Wireless Virtualization Framework,” MONAMI’10.

[ZEN2012] Hongyi Zeng, Peyman Kazemian, George Varghese, and Nick McKeown. 2012. Automatic test packet generation. In Proceedings of the 8th international conference on Emerging networking experiments and technologies (CoNEXT '12).