



# **NetWorld 2020 ETP**

**Expert Working Group on**

**Network and Service Virtualization**

**White Paper**

**Chair: Arturo Azcorra, IMDEA Networks  
and U. Carlos III**

[arturo.azcorra@imdea.org](mailto:arturo.azcorra@imdea.org)

## List of Contributors

Contributors	Company/institute	e.mail address
Arturo Azcorra	IMDEA Networks and U. Carlos III, Spain	<a href="mailto:arturo.azcorra@imdea.org">arturo.azcorra@imdea.org</a>
Carlos Jesús Bernardos	IMDEA Networks, Spain	<a href="mailto:cjbc@it.uc3m.es">cjbc@it.uc3m.es</a>
Antonio de la Oliva	U. Carlos III de Madrid, Spain	<a href="mailto:aoliva@it.uc3m.es">aoliva@it.uc3m.es</a>
Xavier Perez-Costa	NEC Europe	<a href="mailto:Xavier.Perez-Costa@neclab.eu">Xavier.Perez-Costa@neclab.eu</a>
Noel Crespi	Institut Telecom, Telecom SudParis, France	<a href="mailto:Noel.crespi@institut-telecom.fr">Noel.crespi@institut-telecom.fr</a>
Barry Evans	Surrey University, UK	<a href="mailto:B.Evans@surrey.ac.uk">B.Evans@surrey.ac.uk</a>
Marcos Álvarez-Díaz	Gradiant, Spain	<a href="mailto:malvarez@gradient.org">malvarez@gradient.org</a>
Holger Karl	U. Paderborn, Germany	<a href="mailto:holger.karl@upb.de">holger.karl@upb.de</a>
Rahim Tafazolli	Surrey University, UK	<a href="mailto:r.tafazolli@surrey.ac.uk">r.tafazolli@surrey.ac.uk</a>
Diego Lopez	Telefonica, Spain	<a href="mailto:diego@tid.es">diego@tid.es</a>
Marius Corici	Fraunhofer, Germany	<a href="mailto:marius-iulian.corici@fokus.fraunhofer.de">marius-iulian.corici@fokus.fraunhofer.de</a>
Thomas Magedanz	Fraunhofer, Germany	<a href="mailto:thomas.magedanz@fokus.fraunhofer.de">thomas.magedanz@fokus.fraunhofer.de</a>
Eduard Escalona	i2cat, Spain	<a href="mailto:Eduard.escalona@i2cat.net">Eduard.escalona@i2cat.net</a>
Joan Antoni Garcia	i2cat, Spain	<a href="mailto:joan.antoni.garcia@i2cat.net">joan.antoni.garcia@i2cat.net</a>
Artur Hecker	Huawei, Germany	<a href="mailto:Artur.Hecker@huawei.com">Artur.Hecker@huawei.com</a>
Mayur Channegowda	University of Bristol, UK	<a href="mailto:Mayur.Channegowda@bristol.ac.uk">Mayur.Channegowda@bristol.ac.uk</a>
Dimitra Simeonidou	University of Bristol, UK	<a href="mailto:dimitra.simeonidou@bristol.ac.uk">dimitra.simeonidou@bristol.ac.uk</a>
David Lund	HW Communications, UK	<a href="mailto:dlund@hwcomms.com">dlund@hwcomms.com</a>
Hamed Al-Raweshidy	Brunel University, UK	<a href="mailto:Hamed.al-raweshidy@brunel.ac.uk">Hamed.al-raweshidy@brunel.ac.uk</a>
Geir Millstein	Telenor, Norway	<a href="mailto:geir.millstein@telenor.com">geir.millstein@telenor.com</a>

## List of Acronyms

5G	5 <sup>Th</sup> Generation
API	Application Programming Interface
CAPEX	Capital Expenditure
HD	High Definition
IT	Information Technology
M2M	Machine to Machine
NFV	Network Function Virtualisation
OPEX	Operational Expenditure
RAN	Radio Access Network
SDN	Software Defined Network
SLA	Service Level Agreement

## Table of Contents

List of Contributors .....	2
List of Acronyms .....	3
Table of Contents .....	4
Executive Summary .....	5
1 Rationale .....	6
2 Research priorities .....	7
3 5G Evolution Roadmap .....	12
4 Summary .....	14
5 Recommendations .....	14

## Executive Summary

This document summarizes the work of the networking community related with network and service virtualization aspects. These aspects are considered as key components of the design of the future networks, and therefore will be important pieces of the research space in 5G networks. Existing Network Function Virtualization (NFV) and Software Defined Networking (SDN) tools represent today a first step on this direction. However, future 5G networks will be much more complex, and will have to cope with much exigent demands, in terms of traffic, service heterogeneity and network capillarity.

This whitepaper first presents the rationale for pursuing the development of advanced network and service virtualization technologies, to then identify and describe some research priorities within that area. Amongst the priorities identified as critical, one should highlight: efficient RAN sharing for multi-tenancy, personalized “follow me” context aware networks, smart orchestration and use of network analytics, co-existence with existing network deployments, “on the fly” virtualization and adaptability, orchestration of different management planes, terminal virtualization and security.

Finally, a 5G evolution roadmap is provided with summary and conclusions, as well as recommendations for further research.

## 1 Rationale

The telecommunications sector is experiencing a major revolution that will seriously impact the way networks and services are designed, deployed and operated in the near future. We have witnessed the evolution from the “voice era” to the “data era”, and now we are just experiencing the mobile data explosion (see

Figure 1). Users consume data from many different devices, which are getting more and more powerful, and from many different locations. Besides, existing communication networks now provide connectivity to diverse services with disparate requirements, including for example machine-to-machine (M2M) applications, which are expected to have an increasingly important role in the next years.




Figure 1: Mobile data growth forecast

Telecom infrastructures have been traditionally based on a complex set of interconnected proprietary hardware appliances running different types of distributed protocols. These protocols usually require specialized vendor-specific configuration tools. More often than not, the deployment of a new service required yet another hardware box to be installed and configured. Besides, all these specialized hardware suffer from short lifecycles compared to the fast innovation pace and unexpected customers’ demands. This heavily increases the capital expenditures (CAPEX). Moreover, the operation of the network has also been traditionally very complex, as it involves the configuration of heterogeneous hardware, with different APIs and tools, often requiring human intervention, and thus increasing the operational expenditures (OPEX).

Networks are evolving towards becoming a very dynamic and flexible environment consisting of virtual resources that can be instantiated and released on demand to timely meet customers’ demands or to optimize operator’s internal processes. These virtual functions are interconnected by virtual links that are also set-up dynamically to best serve multiple network services. The existing Network Function Virtualization (NFV) initiative and Software Defined Networking (SDN) tools and concepts represent today a first step in this direction. However, future 5G networks will be much more complex, and will have to cope with much more exigent demands, in terms of data traffic semantics, the granularity of the deployed functions,

resulting system, connectivity requirements, service heterogeneity, system dynamics, QoS and network capillarity. An efficient control and management of such networks becomes crucial to overcome this increasing complexity, where dynamic and automated functions will play a key role.

In order to advance in this unstoppable trend, there are some areas that need to be further investigated and developed with network and service virtualization presenting many new challenges to the entire networking and IT industries.

In order to tackle these serious issues, operators started to make use of network and service virtualization techniques. Applied to networking, virtualization enables, for example, co-locating multiple instances of network functions and services on “general purpose” hardware, which is shared among users. This first step in the direction of a tighter integration of the networking and IT worlds already brings significant savings in terms of CAPEX and OPEX, as well as other advantages, such as reduced time-to-market to deploy new services, reduced conformity and integration testing, better flexibility to scale and evolve existing services, and lower risk barriers to perform service innovation. Virtualization in networking also enables the creation of virtual networks, i.e., of network views independent of the actual hardware and its topology. Both other topologies and other protocol stacks can be created. And we can even go beyond the current slicing of networking technologies (e.g., VLANs) and also fuse network links, nodes and whole segments to single entities, therefore reducing management efforts.

Moreover, virtualization enables the emergence of new business models, allowing the provisioning of tailored network services for specific tenants and application purposes. Virtualisation brings the opportunity to deploy almost independent virtual infrastructures for each tenant, configured. Through this, network services are efficiently offered to a specific set of communication devices, enabling an easy uniform administration of a domain e.g. an enterprise or an M2M specific domain. Through this means, the virtualization enables the separation of concerns between different specialized communication systems, thus enabling their easy control, adaptation and administration.

However, sharing common physical infrastructures arises security concerns that have to be addressed to ensure isolation among virtual resources and services.

## 2 Research priorities

### **Efficient RAN sharing for multi-tenancy**

Future 5G networks will be denser to achieve the capacity increase provided by the deployment of smaller cells. In this scenario, reducing the overall costs is of paramount importance. One way of doing that is by enabling an efficient sharing of the network infrastructure. Existing sharing mechanisms are limited and not sufficiently transparent to the network tenants. RAN virtualization techniques are ultimately to provide logical isolated pieces (slices) of the access infrastructure to individual tenants, so they can operate them as if each virtual slice were a single real physical infrastructure. A well-coordinated sharing of the access infrastructure yields a higher throughput per area, but there is still a lot of research work needed to achieve that.

Cloud-based RAN centralization techniques help reducing the costs associated with highly dense networks by offloading the intensive processing to a cloud. These clouds will support different deployment architectures, from small and localized clouds in the street cabinet or even on current large-scale antenna poles to data-centres. The advantage of these solutions is that software can be upgraded and scaled on the fly. In addition, the centralized architecture is ideally suited for joint scheduling of radio resources, handling inter-cell interference and implementation of advanced architectural concepts such as carrier aggregation. However, the high capacity and strict requirements on delay and jitter pose a serious challenge to the design of an efficient fronthaul in the centralized architecture. This challenge becomes even more important with the virtualization added on top of the centralized architecture.

Secondly, the sharing of the access infrastructure is not limited to the radio access interface and should be extended to include the backhaul network. Again, mechanisms to properly share the backhaul and offer a virtual RAN+backhaul network are still at very early stages.

Another important aspect to be considered when creating virtual networks on top of an access network composed of heterogeneous radios and very densely deployed small cells is energy efficiency. By doing a proper monitoring and orchestration of the physical resources, base stations can be selectively turned off and on, as required to cope with the different demands. An operator managing the physical infrastructure can dynamically update the configuration of the different virtual networks, to better meet the service level agreements (SLAs) of the different tenants.

The basic concept of cloud-RAN is to separate the digital baseband processing units (BBUs) of conventional cell sites, and move the BBUs to the “cloud” (BBU pool or BBU hostelling) for centralized signal processing and management. Hence, multiple questions must be answered in order to deal with these new concepts, for example, for the different scenarios of the C-RAN architectures, what is the optimum position of the base band pool and where can be placed? Is it near the central office, or close to cell sites? And what is the optimum hybrid wireless –optical fibre can be deployed to minimize the latency and increase the bandwidth?

As conclusion, with RAN virtualization it becomes possible to have new operational models based on sharing of RAN resources. This approach will reduce CAPEX and OPEX costs while improving the overall RAN utilization.

#### **Personalized “follow me” context-aware networks**

Advanced network and service virtualization techniques will allow to dynamically create and move personal networks, that can “follow” the user with minimal or no impact on its network access experience. These virtual networks can be built on top of a variety of network elements at the edge, for example on base stations owned by an operator, WiFi access points owned by the user or even different types of user devices. Such personalized networks can adapt their characteristics to the user profile, location and general context. They also allow providing a homogeneous and customized network access environment to users, regardless of their location and adapting to the end-user device capabilities. By doing this, the users or their devices are no longer concerned with different access networks, access-specific authentication mechanisms, etc. While a similar connectivity experience could be achieved by adding intelligence to the end-user device, the use of context-aware personalized virtual networks removes any dependencies on the user terminal, thus leading to faster deployments.

This faster deployment of new network services on demand will likely benefit from software orchestration mechanisms (e.g., SDN), as they improve the manageability and adaptability of the network, by separating the user and control plane functions and centralizing the latter beyond what can be done today. With approaches such SDN, low level control and management functions can be more easily centralized.

The concept of personalized networks is not limited to the edge, since it can be applied in a more general use case to serve different heterogeneous services on top of the same physical infrastructure. For example, IoT/M2M type of traffic is quite different from HD video streaming. Different virtual networks/domains across the whole network might be set up and configured to provide connectivity to different services. Future 5G networks are expected to handle a large variety of traffic with disparate requirements. Therefore, the orchestration of the network resources adapted to both the context and the requirements of the service is an important topic for research.

#### **Smart orchestration and use of network analytics**

The expected complexity of 5G networks demands some “intelligence” in the network. A logically centralized process can analyse the network context so as to, for example, route the data according to the current status or even depending on expected/predicted future events. Additionally, this process can also have self-

awareness and learning capabilities for learning from the consequences of its past actions. Cognitive Networking in 5G should be an inherent characteristic of the SDN. The process will provide automatic adaptation policies to the different network areas.

Current “network monitoring” techniques are limited, as they focus on just portions of information, without taking a look at the whole “network picture”, across domain boundaries even into capillary networks. These limitations can be overcome using network analytics tools as an additional source of “network monitoring” information, by exploiting advanced big data analysis techniques. Rather than just providing status updates the monitoring analytics will customize measurements according to the need e.g. if device X utilization is beyond 80% for more than T time in a particular user network then apply load balancing policy based on user profile.

Since the instantiated and released functions are believed to become very general in the future, the orchestration should be dynamic enough to be able to *a)* keep track of all available and used resources, *b)* determine the best available resource containers for any required function within delays adapted to the function semantics, *c)* implement different scheduling strategies to best adapt to the different context and different natures of resources (storage is different than computing, etc), *d)* deallocate the unused functions and release resources at the end of the execution/lifecycle, *e)* dynamically uphold the required SLA by elastically adapting the allocated functions through additions, re-allocations, etc. Additionally, intelligent orchestration can have objectives like increasing the infrastructure utilization ratio or the number of simultaneously fulfilled user reality instantiations, etc. The orchestration mechanisms per se can be centralized or distributed, depending on the precise requirements on the orchestration. Centralized mechanisms will be able to achieve higher accuracy at the expense of higher requirements on the orchestration platforms or lower overall performs (in terms of request throughput or delay); distributed approaches will suffer from temporarily incoherent views, therefore requiring operations on partial, local views.

### **Co-existence with existing network deployments**

A seamless integration of existing networks with future virtual ones is a critical system design aspect that needs to be carefully tackled. Multiple vendor, technology and administrative domain feed to the problem of integration and a future proof solution is required. Network operators cannot update all their deployments towards a fully virtualized environment, but they will start by doing some green field networks that will co-exist with legacy systems. This requires standardized interfaces and the definition of proper hooks to link the virtual network entities with the real ones. In order to be able to function gracefully, the virtualized environment has to be integrated with the physical environment such as Remote Radio Heads, Internet points of presence, interconnection between different tenant domains or integration with transport nodes. It is outmost important that the end-to-end communication characteristics are supported in an integrated manner over all these domains.

Virtualisation is also a key enabler for the efficient deployment of network entities on top of different tenant infrastructures, types such as the satellite, fixed and mobile terrestrial systems via a unified and virtualised management infrastructure for both systems.

In this context, it is especially important to consider co-existence aspects with 3GPP network architectures, allowing parts of a mobile operator’s deployment to be partially or fully virtualized while keeping other parts of existing deployments.

### **“On the fly” virtualization and adaptability**

In order to quickly and flexibly adapt to all the different requirements and constraints imposed by the heterogeneous services that future 5G networks are expected to serve, network virtualization techniques need to be much more agile than current technologies, in setting up new virtual network entities with programmable data forwarding path whilst keeping an optimal energy consumption balance. This feature is

especially required in order to obtain a fast-adaptable robust network addressing features such as load balancing, high availability and redundancy as expected from a carrier-grade operator system. From the booting time required to have a new network entity up and running, to temporal addition/removal of networks, to the time needed to move a virtual entity from one physical location to another one, all these procedures need to be intensively researched and optimized. Alternatively, other abstractions than using Virtual Machines should be considered. There are, for instance, several promising paths within IT security and server/host virtualization communities that allow extremely rapid creation of isolated execution containers without the overhead of booting a new system.

At the same time, for a full-fledged network virtualization, server and host virtualization are insufficient alone. To achieve independent, isolated virtual topology creation and ubiquitous virtual attachment, both link/path and node aggregation and splitting techniques will require both new protocols and processing rules in the routers and other network devices.

### **APIs and SLAs to external actors**

Network virtualisation enables creation and co-existence of multiple networks over the same physical infrastructure each serving different demands and/or services. In such scenario, it is important to enable external actors, such as service and application providers, to use appropriate configuration interfaces. As the complexity of each virtual environment is foreseen to be equivalent to a physical deployment with a high level of scalability, it is foreseen that a simplified governance API will be provided, for each tenant in which full control is provided however lacking the automatic adaptation features. For example, a content distributor may wish to configure a virtual network in a way that its video delivery service is assured to meet the customer SLA. This is one of the key technical enablers towards the emergence of new business models in the ICT sector in the coming years.

### **Orchestration of different management and control planes**

Several approaches towards network virtualization have been researched during the last years. Some of them have been brought to reality in the form of specialized products. One common characteristic to all of these initiatives is the lack of a unified network management and control protocols to manage and configure the different virtual appliances.

In the future virtualization-based networks, advanced management techniques are required to manage the deployment and to follow the dynamic nature of the network and their integration into the end-to-end communication infrastructure. Following the path initiated by the virtualization platforms in the computer world, mechanisms for the group management of network appliances, the dynamic creation and setup of network functions and their associated entities are important research issues. Management and orchestration also requires awareness of the underlying resource availability, while pure processing and storage tasks for network functions can be shifted to parts of the network where sufficient resources are available, the actual underlying link capacity will be limited by the physical boundaries of the transport medium. Management functionality needs to be resource aware and to consider not only the most opportune location to execute virtualised network functions but also the availability of the underlying resources (e.g. honouring the difference between moving virtual functions for execution into capillary networks or into a cellular access network).

### **Terminal virtualization**

Network virtualization must not be limited to the infrastructure providing connectivity, but should extend to the end-user device. A customer might be subscribed to multiple services and access networks simultaneously enabled by virtualization techniques at the mobile terminal.

It would be useful to provide different isolated services using the same terminal. For example, a user might get access to enterprise services via one virtualized end-user terminal, while accessing personal multimedia content using a different virtualized device, optimized to receive that kind of content. This trend is already

discernable at the horizon with the current trends for BYOD and the resulting problems in the enterprises and the recent introduction of isolation functions in the modern terminals.

### **Service delivery management**

Due to the new business possibilities that network virtualization offers, an efficient service delivery management has to be carefully implemented. Relations between resources, services, providers and users can bring a complexity that virtualization may hinder. Different service provider roles may involve hierarchy of virtualized infrastructure increasing the complexity of service deployment. Therefore, a proper service management should consider the required information (e.g., ownership, access rights, delegation, context, SLAs) to allow a multi level management while keeping consistency among the different layers.

### **Security**

Security is a cross-cutting topic that needs to be carefully considered by design from the start. Secure mechanisms must be developed that allows access to only authorized parties (both in human form or through agents) in creating/configuring a virtual network. It must be flexible enough to accommodate the increased, and multi-dimensional dynamics introduced by fine-grained virtualisation. Additionally, isolation between virtualized networks has to be guaranteed as well as providing trust for the virtual infrastructure users on the virtual service enablement.

Current networks are comprised of secure pipes that terminate in insecure endpoints. Virtualised endpoints, agents or network functions will become more granular, increasing the number of secure pipes and the complexities of trust, overcomplicating modern day public key schemes. New techniques such as Identity Based Encryption (IBE), Attribute Based Encryption (ABE), functional encryption and fully homomorphic encryption offer new solutions, which should be addressed from the start. New key management techniques are of primary importance.

### **Performance and QoS**

The virtualization (possibly iterated) of a network and ICT infrastructure can bring several challenges in terms of the level of QoS that can be guaranteed in and by the virtualized resources at each level. The performance behaviours (in terms of, e.g., bandwidth, delay, jitter) of virtualized network resources (e.g. nodes, links) can be severely disrupted if proper QoS isolation techniques are not applied when implementing virtualization. This applies to both past/present data plane virtualization techniques (e.g. NFV), and future ones. Also, the challenge is made more complex by the nature (e.g. deterministic vs statistical multiplexing) and heterogeneity (virtual networks made of homogeneous vs heterogeneous technologies and resources) of the physical data plane to be virtualized.

### **Cost and Performance Analysis**

It is important to study the impact of virtualizing the network functions on the key metrics cost and performance. This will in turn be the key input to the strategic question of which elements to virtualize in the network and how. To this end sub-variables such as elasticity, link utilization, portability and *processing overhead* need to be identified and evaluated which will influence the cost and performance. For example, it needs to be investigated as to how will the virtualisation of a network function and the use of a hypervisor impact the *processing speed* compared to a physical network function. The key metrics should be evaluated depending on the actual use cases that is studied.

### 3 5G Evolution Roadmap

Timeline Technology	< 2015 Features	<2020	2020+ Features
NFV – Network Function Virtualisation	<ul style="list-style-type: none"> <li>• Serving Gateway</li> <li>• Packet Data Network Gateway</li> <li>• CloudEPC</li> <li>• Mobility Management</li> <li>• Pre-defined function migration</li> </ul>	<ul style="list-style-type: none"> <li>• Billing as a Service</li> <li>• Terminal virtualization</li> <li>• Radio and capillary network functions virtualised</li> <li>• On demand function migration</li> </ul>	<ul style="list-style-type: none"> <li>• Virtual firewalls</li> <li>• Automated function migration</li> </ul>
SDN – Software Defined Networking	Centralized Network service orchestration	<ul style="list-style-type: none"> <li>• Automated Flow Management</li> <li>• Automated Life Cycle Management</li> <li>• Distributed Network service orchestration</li> </ul>	<ul style="list-style-type: none"> <li>• Policy driven self-organization and management</li> <li>• Automated placement of virtual firewalls</li> <li>• Automated management and monitoring functions (performance/reliability/SLA management/ security)</li> </ul>
NaaS – Network as a Service PaaS / IaaS	VPN over virtualized networks	<ul style="list-style-type: none"> <li>• APIs for Management and Data Planes</li> <li>• Automated network analytics</li> </ul>	Freely definable private or open networks
Hypervisor	Optimised for OS or specific hardware platform	Bare metal hypervisor with increased hardware support	Dedicated designed hypervisor hardware
Security Trust	Centralised trust anchor	Distributed trust	Fully Dynamic Trust
Core	<ul style="list-style-type: none"> <li>• Some network functions provided in virtualized environments (NFV)</li> <li>• Initial deployment of SDN-based transport networks</li> <li>• Isolation of traffic based on VLAN-like tagging</li> </ul>	<ul style="list-style-type: none"> <li>• Definition of all-virtual core networks (e.g., vEPC)</li> <li>• Enablement of automatic and dynamic connection of new RAN &amp; Core elements to the operator network</li> <li>• Smart monitoring and</li> </ul>	<ul style="list-style-type: none"> <li>• On-demand function specific and context-aware core deployment using shared transport network</li> </ul>

	<ul style="list-style-type: none"> <li>• Monitoring of overall status</li> <li>• One core fits all approach</li> </ul>	<p>trend analysis based on Data Analytics</p> <ul style="list-style-type: none"> <li>• Definition of new function specific customized core functions</li> </ul>	
Access	<ul style="list-style-type: none"> <li>• Basic RAN sharing based on eNodeB location sharing among operators</li> <li>• Static configuration of RAN splitting among operators</li> <li>• Decommission of points of attachment based on static rules</li> </ul>	<ul style="list-style-type: none"> <li>• Generic virtualization of heterogeneous RAN elements</li> <li>• SDN at the Wireless Interface</li> <li>• SDN-based control and configuration</li> <li>• Interfaces for the dynamic creation and decommission of virtual points of attachment</li> </ul>	<ul style="list-style-type: none"> <li>• On-demand Heterogeneous generalized RAN sharing</li> <li>• On-demand deployment of customized RAN elements</li> </ul>
Multi-tenancy and Orchestration	Static agreements between operators based on SLAs	<ul style="list-style-type: none"> <li>• Definition of APIs for the coordination of different virtualized operator's cores over the shared RAN and transport infrastructure</li> <li>• Unification of network management control operating over shared infrastructure</li> </ul>	On the fly sharing of network infrastructure through dynamic orchestration of network functions and elements belonging to different operators

## **4 Summary**

Service and network virtualization techniques are key enablers for future 5G networks. By evolving current network infrastructures into highly dynamic and flexible virtual environments, operators would be able to cope with dynamics of mobile users' demands whilst keeping control over CAPEX and OPEX. Addressing the important research challenges raised by otherwise promising network virtualization technologies will ensure sustainability and growth of mobile broadband networks beyond 2020.

## **5 Recommendations**

The following are recommendations for research concerning service and network virtualization.

- R1) Initiate activity on advanced network and service virtualization to enable efficient RAN & backhaul sharing as well as efficient integration of satellite and terrestrial domains.
- R2) Develop technologies that, based on specific services/users/networks contexts, allow dynamic and flexible creation and operational control of both of virtual networks and the underlying infrastructure resource containers
- R3) Prioritize network research on advanced network virtualization.
- R4) Design mechanisms to ensure co-existence of virtual networks with existing infrastructure, security and efficient management of virtualized networks and services.
- R5) Investigate the impact of virtualizing the network functions on cost and performance.