Future networks – Agile architecture

Network virtualisation

Introduction

Electronic communications networks are composed of a great many separate elements, each with its own specific function: routers forward traffic, firewalls create a barrier between the different parts of the network, etc. In the past, each of these functions was performed by a different piece of hardware. The concept of virtualisation, which consists of separating hardware from software, has been tried and tested in classic computing, and today is driving two revolutions in the telecoms universe:

- The ability to decouple hardware from software in network equipment: several network functions can, for instance, be run independently using the same generic hardware. Network functions can also be moved from one piece of hardware to another, a process referred to as Network Functions Virtualisation or NFV.
- The ability to configure network equipment on the fly according to an application's or service's needs using a "network controller". This is referred to as Software Defined Networking or SDN.

The purpose of NFV is therefore to make physical network equipment multifunctional, and so enable that hardware to perform a wider array of tasks – with each function becoming a software programme, rather than requiring its own piece of hardware. The purpose of SDN is to make traffic routing and processing programmable. These are two separate concepts but developing side by side, and the availability of virtualised network functions provides increased flexibility for configuring and orchestrating networks.

It is easy to imagine the resulting gains in management efficiency that network operators will derive from adopting these two concepts. They can also expect to reduce fixed costs by acquiring generic hardware to run software functions, instead of dedicated hardware. These potential gains nevertheless need to be tempered with the addition of new recurring or one-time costs: the cost of acquiring and integrating licensed software, and retraining the staff that will be operating these new technologies.

It is also clear that these technologies will have an impact on operators' business models – e.g. in terms of cost by reducing capital expenditures (Capex) and increasing operating expenses (Opex) – and on telecoms equipment suppliers' business models – with a parallel increase in the proportion of recurring revenue, and by encouraging them to roll out new services or, on the contrary, by forcing them to share their traditional markets with other IT industry players. This shift could also open the way for new products: an operator could sell third parties (other operators or business customers) a service for hosting network functions on its own network, or even a complete, turnkey virtualised network, including the core network. The use of software-defined technologies and components could also usher in new players along the entire value chain (new operators, new equipment suppliers, new kinds of software pure player, etc.) and drive veteran players to alter their position along the chain. The advent of these technologies creates a number of challenges, and raises questions over how networks today are designed, operated and regulated:

- How to reap the full potential of these technologies for the sake of innovation and competition, while also protecting national industrial capacity?

- How to ensure a high enough level of security when historically disjointed functions now run off the same piece of equipment?
- How to ensure compliance with the principles of net neutrality on a network that can be reconfigured in virtually infinite ways in real time?
- Is current regulation (and especially operators' legal obligations) future proof or, on the contrary, does it need to be adapted to take these developments into account?
- Are current standardisation efforts enough to guarantee that the incorporation of these new technologies does not result in a degree of technological lock-in?

This brief, which is the fruit of a first round of interviews, will explore these questions as a way to identify those issues that will then be the subject of more in-depth analysis.

1 Network functions virtualisation (NFV) software-defined networking (SDN)

1.1 Description of the technologies and the possibilities they create

1.1.1 Network functions virtualisation (NFV): making hardware multifunctional

An electronic communications network is composed of several elements, each of which has a very specific function, which includes the equipment in charge of controlling access to the network, firewalls, routers, gateways that provide an interface between two separate areas, service platforms, databases, etc.[1] Generally speaking, up until now, these functions have been intrinsically linked with the hardware on which they run – so equipment suppliers typically sell the hardware and its function as a single product. Network Function Virtualisation or NFV breaks up this couple by using mature solutions that originated in the world of IT when cloud computing began to catch on. We can draw an analogy with smartphones which perform the functions of several device: a single object can serve different functions, such as a phone, a camera, a game console and a pedometer, because each is performed by software[2]. By decoupling the equipment's functions (software) from its hardware, virtualisation makes it possible to acquire this software independently, and install it on standard servers. These high-capacity servers are spread out over several points of presence owned by an operator[3], i.e. datacentres.

The NFV concept was introduced in 2012 in a White Paper co-signed by 13 operators, summarising the advantages of virtualisation, and calling on the industry to develop it for telecoms operators to use. Through a dedicated working group, the European Telecommunications Standards Institute (ETSI) published a first set of common specifications for implementing network virtualisation (cf. Annex 1 for more details on the structure of NFV) in a multi-vendor environment[4].

---

[1] In current non-virtualised networks, a single piece of hardware can support several different functions: e.g. CPE (Customer-Premises Equipment) is a router located on the user's premises (e.g. a company's offices) which also serves as a firewall and performs quality of service functions. Their configuration nevertheless remains specific, rigid and hard to change on the fly.

[2] In some cases, it has nevertheless been necessary to add new components to the smartphone to enable it to perform this function, e.g. an image sensor.

[3] A typical architecture generally includes a large main site (acting as the central hub) interconnected with several more remote secondary locations over superfast access lines.

[4] The purpose of adopting a common framework is to eliminate the specificities of each supplier's model, and allow the users of the technology (in this case telcos) to operate a single platform whose components are supplied by several vendors.

1.1.2 Software defined networking (SDN): programming traffic routing

A network's architecture is composed of a set of links and nodes that are interconnected through distributed routers. To relay traffic streams between two points on the network, routers perform two distinct functions: signalling (activating routing algorithms) and transport (forwarding traffic).

- The purpose of the routing function is to calculate, select, establish and maintain the path(s) or route(s) capable of relaying traffic between the network's different points. It requires substantial processing power to establish the network's topology and to calculate the optimal route at any given moment based on cost criteria (called metrics) such as minimising the number of routers employed and latency (travel time).
- The purpose of the forwarding function is to relay traffic streams efficiently by steering the data packets received at each network node via the router's incoming network interfaces to the router's appropriate outgoing network interfaces, according to a routing table, powered by the routing function.

Alongside the development of network virtualisation we are seeing the development of Software-Defined Networking (SDN) which makes it possible to relax, or eliminate, these constraints by programming and customising traffic routing rules using software (cf. Annex 2 for more details). Signalling (intelligence) is thus separated from traffic routing. By centralising the network's operation (i.e. calculating paths, readjusting them, examining specific processing rules based on service requirements) in the processing intelligence – called the SDN controller which, logically, is centralised – and by distributing intelligence (notably knowledge of the surrounding topology), it becomes possible to deploy an infrastructure composed of routers that essentially become traffic directors. And SDN network's paradigm is thus akin to the concept of overlay networks.

Because an SDN is centralised, network administrators are able to respond quickly to changing needs, thanks to the (decentralised) information they receive, by streamlining resource allocation according to demand. They are also able to reorganise the network more easily, e.g. to create virtual private networks (VPN) or to automatically alter forwarding rules according to the service being used. This centralisation of decision-making, coupled with the decentralisation of network intelligence makes the network's design and operation more flexible and, ultimately, less costly.

1.2 Towards a "softwarisatisation" of telecom's traditional technical model

1.2.1 Combining SDN and NFV would enable a cross-fertilisation of the two technologies

NFV and SDN technologies have developed in a parallel fashion, and are part of the same overall transformation of the telecommunications industry, which draws its inspiration from similar developments in the IT industry. The purpose of NFV is to make hardware multifunctional, and so enable it to perform a wider array of functions – with each function becoming a software programme, rather than its own piece of hardware. The purpose of SDN is to make traffic routing and processing programmable. Both technologies thus help deliver increased operational flexibility and agility, making it even easier for operators to expand their business well beyond simply providing connectivity, by developing a range of services. In addition to the specific contributions from each technology, combining the two allows each to enhance the other. An efficient implementation of NFV supposes the automatic and flexible creation of virtual networks linking the different virtualised functions, which can be supplied by an SDN controller[5]. Conversely, establishing traffic routing rules through an SDN controller requires network functions which, if performed in a virtualised fashion,

[5] This is a path explored by ETSI in particular. Cf. ETSI GS NFV-EVE 005 V1.1.1 (2015-12): Network Functions Virtualisation (NFV) Ecosystem; Report on SDN Usage in NFV Architectural Framework.

can result in clear advantages in terms of availability, flexibility and performance under load. It therefore seems likely that the users and developers of the two technologies will work to promote their combined use.

1.2.2 How does virtualisation impact telecoms quality of service?

As developed by ETSI, the concept of NFV defines an architectural framework for thinking about network functions virtualisation, but does not intend to specify the technical means used to implement it.

These virtualised network functions[6] can be implemented using the two main technical procedures for IT virtualisation (detailed in Annex 1), each of which has its own set of pros and cons.

In the one case, a certain latency is introduced, which could prove incompatible with network functions that have strict latency requirements, and so making it impossible to run these functions on the fly.

In the other case, this latency would be affected very little, but volatility-related risks are likely to appear, e.g. more complex traceability, the danger of more distributed malfunctions, difficulty orchestrating the whole, etc.

However, regardless of the virtualisation method chosen, networks will become increasingly reliable as it will be easier for an administrator to replace one function with another under a virtualised model. Thanks to the orchestrator[7], faulty equipment is almost immediately replaced by a different, but in all respects completely identical, virtualised piece of equipment.

2 Ecosystem and value chain

Virtualisation is likely to have a sizeable impact on both equipment suppliers' and telcos' business models.

2.1 Impact on the equipment supplier market

Equipment suppliers can expect three major developments resulting from virtualisation: a change in required skillsets, a change in business model (a shift to variable rather than fixed revenue) and the growing importance that after-sales service is likely to have.

- A lot of the software used for virtualisation is Open Source, and offered as software components that need to be assembled to create a complete piece of virtualised network equipment. The assembly of these components, which are not necessarily designed to work together, requires a special set of skills, however.
- Some equipment suppliers have the expertise needed to perform this assembly, and sell all-in-one licences that make easy to deploy virtualised network equipment on a generic piece of hardware. These licences generally need to be renewed periodically. This approach has a direct impact on equipment suppliers' business model and their revenue profile, including a likely decrease in fixed revenue and an increase in variable revenue.
- Whether as part of a turnkey solution or solutions centred only on the hardware that will serve as the container for the virtualised functions, equipment suppliers can stand out from the competition by offering to integrate the different virtualised equipment themselves, and

[6] Virtual Network Function or VNF.

[7] Cf. Annex 1

by providing (support and maintenance) guarantees, with greater value-added than what they are currently offering with their dedicated equipment, as the fragmentation and multiplicity of virtualised functions make for an increasingly challenging integration process. Some IT industry players already have the needed expertise in this area, and are therefore likely to become rivals for veteran hardware suppliers. Some virtualisation specialists, such as VMware, have developed premium and optimised versions of their software, to capitalise on the advantages of virtualisation (flexibility, resource pooling) while diminishing the inherent drawbacks – such as decreased performance and reliability (cf.1.2.2). Cloud computing companies are also experts in virtualisation and in managing large-scale IT resources, as well as having the network skills needed to relay large volumes of data in little time.

2.2 Impact on telcos

2.2.1 Financial impact

Network virtualisation will probably lead to financial gains over time, although the scale and timeline for these gains is not yet clear, and depends on each individual situation and on the chosen scenarios.

The providers of virtualised solutions are projecting that virtualisation will generate substantial gains. According to an ACG report from 2015[8], virtualisation could enable telcos to reduce their capital and operating expenditures on network equipment by two thirds.

This estimate does seem to be on the high end, however, and these figures need to be put into perspective given that – according to some experts – the virtualised portion of the network represents only 25% of a mobile operator's annual Capex. The other 75% represent primarily spending on cell sites and land.

In addition, if the cost of buying hardware is expected to decreases significantly thanks to virtualisation, the cost of the software used to implement virtualisation is likely to vary considerably. Some operators could integrate Open Source software solutions directly, but those that do not have the skills or resources needed to manage the integration and adaptation of these solutions may prefer to opt for equipment suppliers' licensing solutions, to have access to software that is optimised and easier to integrate.

In both cases, we will see a shift from a fixed cost to a variable cost model, whether due to the need for more qualified staff in charge of handling the integration and adaptation of the software building blocks, for those using Open Source software solutions, or the cost of obtaining licences for those opting for paid proprietary solutions. The question of financial gain comes into play here. Let us take the example of the virtualised router: on top of the usual hardware maintenance costs will be the cost of software licences whose regular, unit price (which may therefore be different from what operators charge) hover around $8,000[9] a year[10]. To compare, the price of an equivalent non-virtualised router is around $35,000[11] and does not require a licence. Depending on the relative lifespan of these two devices, Capex savings could therefore be offset by licence costs. Moreover, in some cases, several virtual machines will need to be deployed to achieve the same level of performance as a dedicated piece of hardware. So medium-term gains will not be as high as those forecast by ACG.

---

[8] Study financed by VMware, "Total Cost of Ownership Study Virtualizing the Mobile Core" https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/vmware-nfv-tco-report-acg.pdf

[9] Source: ITprice.com

[10] Cost of a 10Gbps router: beyond that speed the virtualised equipment would consume too much energy.

[11] Source: http://itprice.com/cisco-gpl/1000v?p=

If one middle path would be to combine the integration of Open Source software and the purchase of licences from equipment suppliers as a way to reduce costs, in practice this approach could potentially mean that operators would lose the guarantees from equipment suppliers. Taking this route could also require operators to separate, or isolate, the virtual machines supplied by equipment vendors from the other virtual machines.

From a broader perspective, the transition to virtualised networks may also involve hidden costs, particularly in terms of training and, under certain configurations, power consumption, even when using paid proprietary solutions.

2.2.2 Operational impact

From an operational standpoint, virtualisation could create three main challenges for operators: the need to retrain their teams, a possible deterioration in their quality of service, and redefining contractual liability when network outages occur.

- In the IT world, a network engineer can administrate a large number of machines (e.g. an average 30,000 machines at Google) whereas, with telcos, engineers today administrate around 100 machines. In addition to the number of machines, their variety is also very different. This means that operators will need to train some of their administrators in virtualised networks. Several experts have pointed out the fact that there is no curriculum in France today that includes this type of training. The virtualisation of telcos' networks will therefore involve a period during which operators will need to get up to speed with this new technology, while continuing to manage their non-virtualised equipment. If telcos do manage to acquire the necessary skills, some may seek to develop virtualised network software themselves, to be less dependent on equipment suppliers: this is the strategy that AT&T[12] has adopted, for instance. Here again, this will require a sizeable investment in training, but will also mean telcos will no longer have to pay for licences, and that they will be able to develop their own functions.
- Virtualisation may also raise questions over the quality of operators' services (in terms of latency, availability or volatility, depending on the technological choices they make — cf. Section 1.2.2). The scale of the negative impact that virtualisation might have on quality of service is nonetheless crucial, and so determine whether it will be minimal, or whether it will affect networks' operational management in a very real way.
- Virtualisation does indeed create a new issue in the area of contractual liability. In the past, for dedicated equipment (switches, firewalls…), the equipment supplier was typically the sole party that was contractually liable when a problem occurred. With the virtualisation of network functions, which become applications running on generic servers, it becomes harder to identify the responsible party when a breakdown occurs: is it the telco, the equipment supplier[13], the integrator or the software provider? Will the guarantees provided by equipment suppliers allow for configurations that combine licensed software and directly integrated Open Source components?

2.2.3 Impact on service provision

[12] "Hitting the Open Road: Software-Accelerating Our Network with Open Source" http://about.att.com/innovationblog/061714hittingtheopen

[13] In 2008, for instance, Danish telco TDC signed a strategic partnership with equipment supplier Ericsson to manage and operate its networks, to accelerate the time to market for 4G. It is not unimaginable that with virtualisation and the impact it has on the business of telecoms operator, this type of partnership will become increasingly common.

Thanks to virtualisation, it will become possible, even on a physical network, to perform network slicing, e.g. to offer products that are tiered by the quality of service provided to users.

Virtualisation also enables operators to share access to their network infrastructure (cf. Section 1.2.2) with third parties (other telcos, users, businesses, verticals, etc.) to host virtualised functions, which themselves can be provided by the telco (as part of a product line) or by the users themselves – such as the machine operation or monitoring functions required by certain industries, or use cases for verticals (connected cars, smart city etc.).

Under a scenario of greater openness, a telco could give its customers the technical ability to operate their own virtual networks on its platform, either by keeping a degree of control over the platform's intelligence (orchestration, supervision, end-to-end view…), or by disengaging completely from the intelligence and confining itself to the role of infrastructure operator. A device manufacturer (e.g. a connected car maker) could thereby manage its own virtual network, and provide the associated services.

Several operators are already planning on monetising access to some of their network equipment.

2.2.4 Impact on competition

Virtualisation could enable new operators to start a business relying on fully virtualised networks, and not having to own any physical infrastructure themselves. These new operators could be both positioned differently (niche or mass market) and operate on different scales (locally, nationally or transnationally).

A new operator could in fact minimise its start-up costs by using cloud solutions, which would reduce its investment costs. In addition, managing virtualised networks requires different equipment and properly trained staff. New entrants, which could hire qualified staff directly, and deploy the most suitable equipment, could therefore enjoy a comparative advantage.

It could therefore become easier to enter the telecommunications market, particularly for outside players that have expertise in virtualisation and the cloud. This in turn could profoundly alter the telecoms market's competition landscape.

Although virtualisation makes it possible to simulate a physical network without having to be its owner, the need to obtain access authorisation from the owners of the physical infrastructure nevertheless remains. The degree of accessibility and flexibility that infrastructure owners provide could prove an impediment to this development. The new products that emerge thanks to virtualisation, especially transnational ones, will also depend on how homogenous access rules are.

3 Identifying the issues and challenges that might require special attention

3.1 Interoperability

SDN and NFV technologies make it possible to design network architectures that function thanks to several independent and customisable software building blocks. These building blocks can be available as Open Source products or developed in-house by some of the players discussed in Section 2. Ideally, these building blocks could be used in most environments, on most infrastructures, and be able to "talk" to each other: i.e. the crucial issue of interoperability.

This interoperability issue can be tackled on two levels: inside a single network, and at the border of different networks (interconnection).

- Within the same network, interoperability between the different components of an SDN and virtualisation-based network architecture helps prevent inefficiencies (such as the need to customise a network application for each virtualisation socket). Ensuring this interoperability

would therefore allow operators to acquire network functions and deploy them on their infrastructure with minimal adjustments, which would reduce the time and cost of getting them up and running.

- At the border between different networks, interoperability will be vital to ensuring the networks' interconnection and so the guarantee of quality of service end to end. One aspect that warrants particularly attention is the interoperability between different SDN areas, notably when it comes to the interoperability of the different orchestrators, which have not yet been standardised.

To be able to programme the network via APIs and orchestrate virtualised functions, the configuration of the equipment and the virtualised network services need to be modelled[14]. The consequences here fall on the network architects who will now need to model these services.

Standardisation work, which has been carried out by collective industry bodies – e.g. within ETSI and ad-hoc organisations/consortia[15] – and the Open Source software community, could help provide a technical response to this interoperability issue. To guarantee this interoperability, the different organisations will need to come together as much as possible on common, or at the very least compatible, standards.

3.2 Third parties' access to physical infrastructures: innovation, competition and protecting industrial capacities

As we saw in Section 2.2.3, virtualisation opens up a range of technical possibilities for third parties on operators' infrastructures and platforms.

The degree to which this technical potential is exploited with ultimately depend on how open the operator is to third parties, which could range from merely marketing a line of virtualised services instantiated by the operator itself, up to giving third parties the ability to design and orchestrate their virtualised functions themselves.

And so arises the question of the rules governing access to physical infrastructure and exposing Application Programming Interfaces (API) to third parties so they can host their virtualised functions and configure their virtualised networks. The API question also gives rise to a competition issue when players favour a given protocol or API in such a way that excludes those they do not use.

More broadly speaking, then, the fact of altering access rules (or not) could have an impact on the competition dynamic and on the fostering of innovation. The question of protecting national players' industrial capacities could also arise.

3.3 Net neutrality and tiered quality of service

As mentioned earlier, virtualisation makes it possible to implement network slicing. Several QoS parameters pertaining to the traffic being relayed can be configured specifically for each slice. The use of SDN in a network centralises and streamlines the process of configuring quality of service for each network, which could make it easier to define specific handling by type of stream.

---

[14] The YANG ("Yet Another Next Generation", created by an IETF working group) data modelling language is an example of a modelling language for these configurations, making it possible to specify a service model (describing its components, its internal interactions, requirements and capacity) that the SDN orchestrator/controller will use, but also with respect to third parties (exposing this service via APIs).

[15] E.g. the TIP (Telecommunication Infrastructure Project) consortium initiated by Facebook, and the not for profit ONF (Open Network Foundation) consortium initiated by telcos, academics and OTT players.

This gives rise to the question of whether these functions comply with the regulatory framework governing net neutrality, and the procedure for enforcing compliance with this obligation.

Here, BEREC has not yet identified[16] network slicing as a violation of net neutrality as such, provided it is used in accordance with the quality of service requirements authorised by Europe's Open Internet Regulation (e.g. for "specialised" services or reasonable traffic management practices)[17]. BEREC has also stipulated that it is up to regulatory authorities to determine, on a case by case basis, whether a service complies with Article 3(5) of the Open Internet Regulation, which stipulates the conditions under which operators can provide optimisation for specific services, with a specific level of quality.

To monitor the quality of service policies implemented on networks that employ SDN, one first path – which would be relatively complex to implement – would be to consult the controller's and the orchestrator's configuration files, and the network event logs. This also raises the broader question of accessing network information to monitor compliance with obligations.

3.4 Security

SDN and NFV technologies give rise to four security challenges: the creation of a single point of failure (SPOF) – which is the result of having a single point of network functions control – requires a guarantee of partitioning between applications, larger attack surfaces, and heterogeneous configurations.

- The centralisation of network functions control constitutes a source of vulnerability for any architecture: if the controller becomes unavailable or fails, it becomes impossible to control the networks that depend on it (configuration, routing policy…). To limit the risks tied to having a single point of control, redundant critical elements can be installed as a way to make the network more robust. Virtualisation makes it easier to deploy this redundant equipment. This is primarily a network robustness and reliability issue.
- The infrastructure sharing made possible by virtualisation allows for different applications to run in parallel on the same physical machine: if security conditions guaranteeing partitioning between the applications running simultaneously are not in place, the effectiveness of shared resources and expected gains could be undermined by the need to run each of the most critical network functions separately.
- The architectures are composed of several, independent software building blocks that may need to talk to each other: the increased number of communication channels may mean an increase in the attack surface, and lead to additional vulnerabilities.
- NFV architectures create a great deal of freedom in how the network's building blocks can be configured. This freedom generates, on the one hand, an increase in the number of potential sources of error and, on the other, a tremendous disparity of configurations, and so making security analysis a more complex affair. The number of parties involved in the network's configuration can also dilute each one's (sense of) responsibility.

Conversely, it is possible that the use of SDN will enable greater control over configurations, and make it easier to change them rapidly. The design of a secure framework to undergird the

[16] A working group devoted to updating BEREC net neutrality guidelines was created for 2019.

[17] "According to BEREC's current understanding and analysis, the Regulation seems to be leaving considerable room for the implementation of 5G technologies, such as network slicing, 5QI and Mobile Edge Computing. To date, BEREC is not aware of any concrete example given by stakeholders where the implementation of 5G technology as such would be impeded by the Regulation." https://berec.europa.eu/eng/document_register/subject_matter/berec/opinions/8317-berec-opinion-for-the-evaluation-of-the-application-of-regulation-eu-20152120-and-the-berec-net-neutrality-guidelines

development of virtualisation is one of the projects being worked on by ANSSI, and identified in the strategic review on cybersecurity published in February 2018.

3.5 Sovereignty

Network virtualisation raises two issues regarding control over national sovereignty: first, operators' obligation to keep certain parts of their business inside the country and, second, the obligation to obtain an authorisation for the operation of certain technical devices and systems.

- Virtualisation enables operators to free themselves of geographical restrictions to perform certain core network, network operation and supervision functions. Some operators in France could therefore choose to run certain parts of their business outside the country, either for reasons of cost, or to pool these operations with similar ones being run by their subsidiaries or sister companies operating in other countries. The regulation nevertheless clearly stipulates that certain activities, notably the deployment and implementation of the means needed to intercept correspondence[18], must take place inside the country. Relocating certain functions abroad, which virtualisation makes easier, can also affect the Government's ability to implement its cyberattack detection capabilities, and the ability to react when a crisis occurs. The fact that functions that have long been performed in-house by operators are now being outsourced to outside vendors can also have an impact on sovereignty when, for instance, these companies are subject to foreign regulation (e.g. the Cloud Act[19]).
- Under current laws[20], technical devices and systems capable of performing interception require an authorisation from the Prime Minister. Once network functions, subject to the above-mentioned authorisation regime, are virtualised, however, the question arises of exactly which elements must obtain this authorisation. Would it be only the virtualised function itself, or would an authorisation also be needed for the cloud infrastructure, the physical hardware and operating systems that could be used to run the virtualised functions for which an authorisation is required?

It is therefore vital that operators consider the regulatory and legal constraints, and how they are likely to change, to be able to establish a clear picture of the exact role that virtualisation will play in their future networks and, in turn, the gains they can reap as a result.

Here again, these questions of sovereignty fall under the purview of the federal government in particular (France's National Cybersecurity Agency, ANSSI, and the General Secretariat for Defence and National Security, SGDSN, among others).

[18] CPCE Article D. 98-7 III

[19] https://www.congress.gov/bill/115th-congress/senate-bill/2383/text

[20] C.f. Article R.226-3 of the Penal Code and the Order of 11 August 2016 amending the Order of 4 July 2012 setting the list of devices and technical systems covered Article 226-3 of the Penal Code.

Annex 1

A history of virtualisation specifications

The concept of NFV was first introduced in 2012 in a White Paper[21] co-signed by 13 operators, summarising the advantages of virtualisation, and calling on the industry to develop it with telecoms operators in mind. Through a dedicated working group, ETSI (European Telecommunications Standards Institute), published a first set of common specifications for implementing network virtualisation in a multi-vendor environment. In its simplified version, the ETSI NFV model is structured around three main constituent parts:

- NFV infrastructure, a generic infrastructure in the form of servers, switches, databases, etc. and which is agnostic with respect to the applications it supports;
- Virtual Network Functions (VNF), which leverage the progress made in IT virtualisation, and which are instantiated on the fly in NFV infrastructure;
- Management and Orchestration (MANO), which controls and manages these VNF (managing lifecycles, controlling their elasticity, choice of the physical server on which to run this or that VNF etc.). The orchestrator steers the VNF's behaviour through a series of dedicated descriptor files supplied by the vendor, containing their properties, instructions describing the hardware resources required to deploy them through the Virtual Infrastructure Manager, or VIM (e.g. internal/external connectivity type and datarate, processing/storage capacity, etc.) and the conditions that could hamper their assembly with other VNF. The orchestrator is, in essence, the guarantor of VNF composition and orchestration by having an end-to-end view of the service.

The first proposals for implementing VNF are based on virtual machines (VM), wherein physical resources (infrastructure) are partitioned, decoupled by a dedicated layer of software, and where each virtual machine embeds the software of the network function to be virtualised. Thanks to this partitioning, several operators can control virtual machines deployed on the same physical machine. These virtual machines have their own operating systems, however, and typically do not enable direct access to the physical resource. As a result, virtual machines are slightly less reactive than the physical machines they emulate. If this slight latency caused by virtualisation is not a problem when using a classic computer, it is incompatible with certain network functions with strict latency requirements, and thus rules out the ability to run these functions on the fly. There are alternative proposals that argue for the use of containers, a miniature variant of the virtual machine that involves a more porous partition between the virtual machine and the physical infrastructure. Containers are more reactive than virtual machines, and they have been used for some time on cloud platforms for providing web services[22] and micro-services (native Cloud), but are also more volatile[23].

[21] *Network Functions Virtualisation: An Introduction, Benefits, Enablers, Challenges & Call for Action*. White Paper published at the *SDN and OpenFlow World Congress* (2012)

[22] These are the component building blocks for applications such as Facebook, Gmail, Amazon apps, etc.

[23] Volatility is tied to a container's granularity. Unlike a virtual machine, the container is a micro-component which creates its own set of risks: traceability is more complex, more distributed risk of failure, more challenging to coordinate and orchestrate the whole, etc.

Annex 2

SDN

To perform their function of traffic directors, routers employ substantial signalling and processing power (notably for finding neighbouring equipment, calculating the optimal route, performing certain specific tasks – such as packet inspection, filtering, load balancing, etc. – as an integral part of standard routing and packet forwarding tasks) and often require static programming of traffic management rules. Software defined networking (SDN) creates the ability to relax, and possibly eliminate, these constraints by programming and customising traffic routing and management rules through software, according to the service's requirements/objectives.

This software-based programming is made possible by decoupling the data transmission/processing layer (i.e. Data Plane) from the control/signalling layer (i.e. the Control Plane) which is typically incorporated into the same equipment. This decoupling was initially formulated by an ITU-T recommendation, then theorised by ONF[24] in 2014 through an architecture of three distinct layers that interact through Application Programming Interfaces (API[25]): the software layer (which includes the SDN network's end users' apps), the control layer (which supplies the centralised control functions that supervises and influences the network's forwarding behaviour, embodied by the SDN controller) and the infrastructure layer (which includes all of the network equipment in charge of performing packet switching and forwarding, as instructed by the SDN controller).

SDN makes it possible to manage the forwarding table from a controller that is naturally centralised, which makes the network's operation and management very flexible. The three-layer architecture with well documented API between the layers creates an independence with respect to suppliers, and so freeing the network operator from being locked into any one vendor's (technical and so business) rationale.

[24] Open Network Foundation (ONF), SDN Architecture (06/2014)

[25] Application Programming Interface (API) which has the role of clearly defined façade enabling a software entity to communicate with another entity. Widely developed in the world of IT to facilitate the creation of apps and software services, APIs are typically made available in the form of software libraries (Open Source or proprietary). Thanks to API abstraction, concerns relating to the technical elements that are specific to the entity's internal operation are masked, and so plays a key role in fostering interoperability.

Annex 3

Bibliography

- ETSI, *Network Functions Virtualisation: Introductory White Paper*
  https://portal.etsi.org/nfv/nfv_white_paper.pdf
  https://www.etsi.org/technologies/nfv
- AT&T, *AT&T Vision Alignment Challenge Technology Survey, AT&T Domain 2.0 Vision White Paper*, November 13, 2013.
- Martin Taylor, *The application of Cloud Native design principles to network function virtualisation*, Metaswitch.
- Bruno Chatras, *La virtualisation des fonctions de réseaux de télécommunication*, Revue Telecom No. 181, May 2017.
- 5G-PPP Software Network Working Group, *From Webscale to Telco, the Cloud Native Journey*, Cloud Native White Paper, July 2018
- European Commission, *Implications of the emerging technologies Software-Defined Networking and Network Function Virtualisation on the future Telecommunications Landscape (SMART 2005/0011)*, 2016

  http://ec.europa.eu/newsroom/dae/document.cfm?doc_id=44557
- BEREC, *Input paper on Potential Regulatory Implications of Software-Defined Networking and Network Functions Virtualisation (BoR (16) 97)*, 2016

  https://berec.europa.eu/eng/document_register/subject_matter/berec/download/0/6088-input-paper-on-potential-regulatory-impl_0.pdf
- IDATE Digiworld Research, *Virtualisation in Telco Networks: which markets for SDN and NFV, and what perspectives with network slicing for 5G?*, September 2017
- *SDN and NFV Simplified: A visual guide to understanding Software Defined Networks and Network Function Virtualisation*", Jim Doherty, 2016 Pearson Education.
- Diego Kreutz, Fernando M. V. Ramos, Paulo Verissimo, Christian Esteve Rothenberg, Siamak Azodolmolky, Steve Uhlig, *Software-Defined Networking: A Comprehensive Survey,* IEEE Surveys & Tutorials on communications.