*eChallenges e-2014 Conference Proceedings*
*Paul Cunningham and Miriam Cunningham (Eds)*
*IIMC International Information Management Corporation, 2014*
*ISBN: 978-1-905824-45-8*

# Offloading Security Applications into the Network

Antonio LIOY[1], Antonio PASTOR[2], Fulvio RISSO[1], Roberto SASSU[1], Adrian L. SHAW[3]
[1]*Politecnico di Torino – DAUIN, c. Duca degli Abruzzi 24, Torino, 10129, Italy*
*Email: antonio.lioy@polito.it, fulvio.risso@polito.it, roberto.sassu@polito.it*
[2]*Telefónica I +D, c. Don Ramon de la Cruz 82, Madrid, 28006, Spain*
*Email: antonio.pastorperales@telefonica.com*
[3]*Hewlett-Packard Laboratories, Long Down Avenue, Bristol, BS34 8QZ, United Kingdom*
*Email: adrian.shaw@hp.com*

**Abstract:** Users currently experience different levels of protection when accessing the Internet via their various personal devices and network connections, due to variable network security conditions and security applications available at each device. The SECURED project addresses these issues by designing an architecture to offload security applications from the end-user devices to a suitable trusted node in the network: the Network Edge Device (NED). Users populate a repository with their security applications and policy, which will then be fetched by the closest NED to protect the user's traffic when he connects to a network. This setting provides uniform protection, independent of the actual user device and network location (e.g. public WiFi hotspot or 3G mobile connection). In other words, a user-centric approach is fostered by this architecture, opposed to the current device- or network-based security schema, with cost and protection benefits and simultaneously enabling new business models for service and network providers.

## 1. Introduction

Nowadays, most people in developed countries exploit several personal devices to access the Internet, ranging from "traditional" terminals, such as desktops, laptops, tablets and smartphones, to new devices such as smart TV, car infotainment systems, and more. Furthermore, this number is expected to increase in the coming years due to smart Internet-of-Things and wearable devices, such as Google Glass [1].

While this scenario empowers the user with new and attractive capabilities, it also, at the same time, poses significant challenges in terms of security, particularly with respect to protection of the devices (and their users) from external threats.

First, many devices have limited computational and storage resources, particularly embedded devices and mobile terminals, and are often further constrained by severe limitations in terms of power consumption. As a consequence, complex protection applications (like malware detection with a large signature database or VPN client with strong encryption) may not be easily executed on all devices, hence exposing their communications to security threats like malware infection or data interception during the network transfer.

A second issue is related to the variety of network connections exploited by end-users that want the ability to connect to the Internet anywhere and anytime: home xDSL, public WiFi hotspots, company networks or 3G/4G mobile connections are some examples. However, this flexibility comes at a price: the intrinsic security level of these various connections is highly different. For example, a user is typically exposed to more threats when connecting from a public hotspot than when connecting from the corporate network as this last one usually includes a sophisticated border firewall.

 

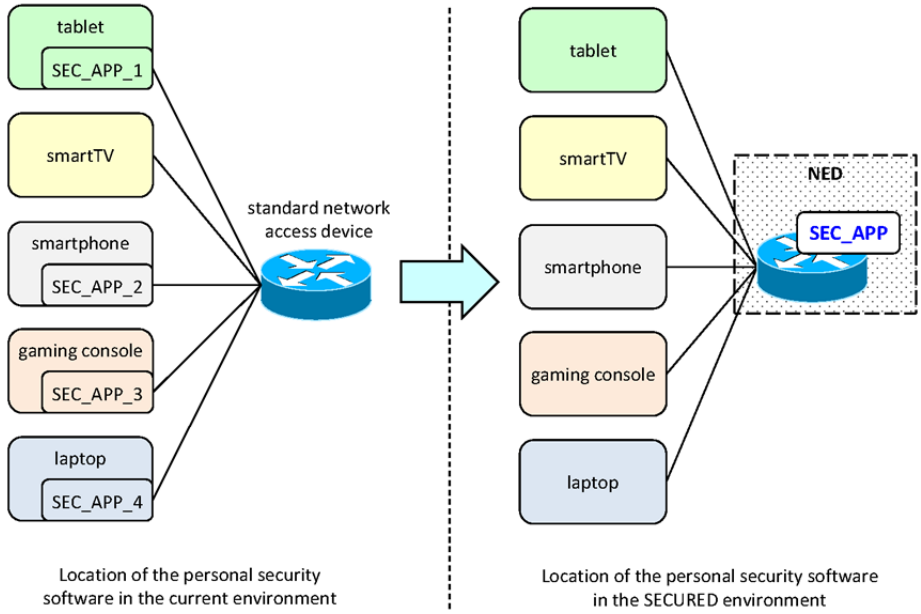


*Figure 1: moving security applications from user terminals to the network.*

Third, given the high variability in hardware and software platforms adopted by user devices, the same applications may not be available for all devices. For example, powerful parental control applications are normally available for desktops and laptops while the same software may not be available when browsing the Internet from other devices, such as a smartphone or a smart TV, hence leaving kids unprotected. Additionally, some devices are not promptly updated with security patches [2] or don't even have the ability to run user-selected security applications, as is the case for most smart-TV sets. This boils down to different security applications being available at the various devices and variable security level of their base software.

As a consequence of this overall scenario, the user is forced to buy different protection applications for his many devices (with the associated problems of cost and different configurations and capabilities) yet might not achieve uniform protection (e.g. in case the application does not exist for the given platform or the user connects from a machine that cannot be secured, such as a public kiosk): a real nightmare for security-conscious citizens.

## 2. Objectives

This paper proposes a possible solution to the above mentioned problems, based on a network application offloading approach (Figure 1).

In a nutshell, we move the security protection from the user terminal to the (closest) *network edge device (NED)*, which can be represented by a properly enhanced access point, switch or router operating at the edge of the network. According to this approach, users will configure the desired security countermeasures (applications and policies) only once and the NEDs will have the responsibility to apply them automatically to all the network connections started by the given user, regardless of the physical terminal in use and the current network attachment point.

The main advantage of this approach consists in transforming the current protection approach from network-based (or device-based) into a new user-centric paradigm, hence delivering personalized protection independent from the user's device and location. In addition, this would no longer require installing specific software on each user device, which simplifies management and, potentially, reduces power consumption, hence offering to devices with limited capabilities the same level of protection of more complex platforms.

Of course this approach requires a good deal of trust in the NED, which is a very critical issue given the recent public disclosure that national security agencies intercept a massive amount of personal communications. This is one of the reasons why we took particular care in defining strong technical solutions (described in the next sections) to formally guarantee high levels of trust in the NED. In any case, if the user does not trust the NED or prefers a double line of protection to implement defence-in-depth – the NED can be used as a supplement to standard device- or network-based security measures. In other words, using the NED does not force the user to give up his legacy security solutions.

## 3. Methodology

In order to fulfil our objectives, we need different logical components.

First, we need to enrich the network with the capability to recognize which user is currently connecting to it. A wide range of options are available, starting with a traditional user/password authentication (maybe through a captive portal), 802.1x authentication, or even more unconventional approaches, such as facial recognition (which requires an additional software on the user terminal) that exploits the fact that most of the modern user terminals include also an integrated webcam. It is worth noting that user identity is not needed to accept or refuse the right to connect to the network (we assume that this problem is managed in a separate and standard way by the network provider). Instead, it is required to retrieve the user profile and configure the NED with the security applications/policies associated to the user itself. Of course, in case the network provider is also the NED provider then it can choose to use the same system for network access control and authentication towards the NED.

The second component, the NED, is a network device that integrates network and computing capabilities and that is reconfigurable at run-time to execute the required security applications based on the users attached to it. This component represents one of the most critical pieces of the architecture and it needs to provide the proper isolation between different users (both with respect to their network traffic and computations). Additionally, the NED must be trusted as the user must be certain that the NED is really executing only the user's applications and not performing any malicious activities on the user's traffic. Finally, all the traffic between the user terminal and the NED must be protected by a secure channel, so that user's data cannot be read or modified by an attacker before the NED applies the required user-defined protection measures.

The third component addresses the necessity to protect the Internet connection when the user attaches to a traditional network, which does not have the network application off-loading capabilities. This requires installing a small monitoring application on the user terminal that is activated each time the device attaches to a new network and checks that the connection is made to a trusted and secure NED. In case this condition is not verified, the above mentioned application will establish a remote connection (e.g. VPN) to a trusted NED and redirect all the traffic of the user terminal to the remote network node, hence guaranteeing the expected level of protection although with a higher latency.

## 4. Technology Description

The resulting architecture exploits several cutting-edge technologies.

First, to meet the trust requirement, NEDs are equipped with the Trusted Platform Module (TPM) cryptographic chip [3], which – when coupled with appropriate management software, for example TrouSerS [4] – permits to securely measure the software state of the platform and to report it along with the platform identity. This enables users to determine if the device they are attached to is a NED (i.e. is part of the secure infrastructure) and then use remote attestation techniques [5, 6] to verify that the software

executed at the NED is that expected by the user and no tampering occurred both in the NED operating system software and the user's applications. This guarantees to the user that the network operator is not performing malicious or unwanted actions on his traffic.

As a NED will process at the same time the traffic of several users, strict isolation is important, both with respect to traffic and computations, and can be provided with several technologies.

SDN (Software Defined Networking) techniques, such as OpenFlow [7], hold the promise to split the network in different partitions, each one assigned to a specific user, and to confine the traffic of each user in his own slice. This ensures that user's packets will always be processed by a NED before they reach other users.

From the computing point of view, we need a distinct execution container for each user, to allow multiple users to run their applications on the same NED in an independent and isolated way. In this respect, various options are available that provide different trade-offs between isolation, generality and overhead. Dedicated or language-specific Virtual Machines (e.g., Java VMs) are light but would require a major rewriting of the applications and may not offer adequate protection, such as memory and CPU isolation between users. Other technologies such as Docker (http://www.docker.com) or Linux Containers (https://linuxcontainers.org) are more general, still lighter than general virtual machines but are not completely isolated. Full-fledged virtual machines (such as those run by standard hypervisors, e.g. KVM, VMware and Xen) provide the strongest isolation and compatibility with legacy applications, but they introduce a notable overhead particularly in terms of CPU and memory consumption, which prevents to consolidate the applications of a massive number of users on a few NEDs.

Whatever is the execution container chosen for network application offloading, this schema is compatible with the service model proposed in ETSI by the Network Functions Virtualisation (NFV) group which supports "*implementing network functions in software that can run on a range of industry standard server hardware, and that can be moved to, or instantiated in, various locations in the network as required, without the need to install new equipment*" [8].

In addition, the adoption of "industry standard" components, such as OpenFlow (for the networking portion) and KVM (for the computation part), enables our solution to be integrated in cloud-oriented platforms (e.g. OpenStack), hence guaranteeing synergies between different services of a network operator, which could exploit the same infrastructure to provide cloud, networking and security services. Moreover, this allows a NED to easily delegate part of its workload to other machines, such as servers operating in a (remote) datacentre, which can guarantee almost unlimited computational power in addition to cost savings.

Since our approach is user-centric, we are also considering the problem to configure in a simple way the security applications, as the average user will not be an expert in network or system management. To this aim, we are also designing a high-level policy language to express generic protection requirements suitable for the general user. Example are "block all malware", "do not permit access to phishing sites" and "block all content inappropriate for children 7-years old". In turn, the policy expressed with this language will be automatically translated to the actual configurations of the user-selected security applications, thus relieving users from the problem of understanding the nitty-gritty details of security configurations. Furthermore, this would permit network operators to select the security applications that are most appropriate for their infrastructure, while still providing the requested level of security to their users.

An additional benefit of a policy-based security configuration is the possibility to apply more than one policy to the same network connection. For example, a parent when buying a smartphone for his child could sign a contract requiring that the connection to the Internet is

always subject to the policies set by himself, with higher priority, in addition to the policies set by his child. In a similar way, connections of employees in a BYOD environment would be subject to the employer's policies in addition to the ones set by the employee himself. Arbitrary composition of policies, with the proper level of priority is one of the characteristics supported from day zero, with a schema that is hierarchical and general: for example all connections could also always be subject to the policy of the network operator, with the highest priority. The NED would enforce all the policies in the correct order and inform the end-user about the final global policy applied to his network connection.

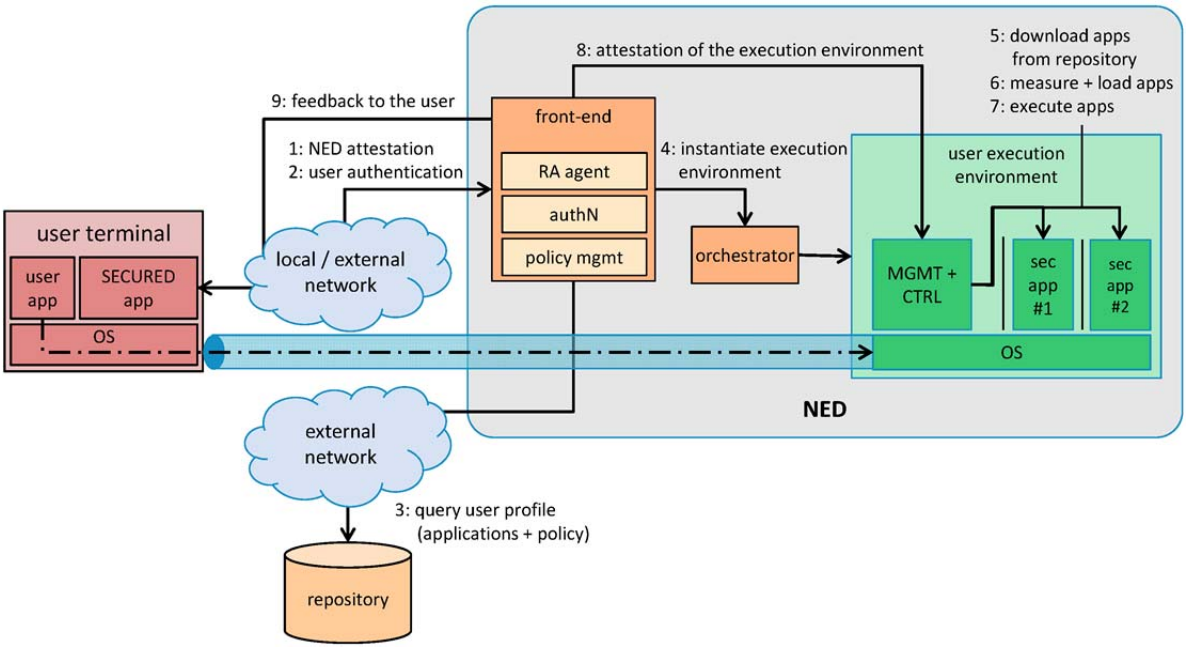


*Figure 2: logical architecture for offloading security applications into the network.*

## 5. Developments

Figure 2 shows the generic logical architecture that we designed to offload security applications to the NED. When it comes to implementation, we foresee the development of two main flavours of the NED: the monolithic and the split ones.

The *monolithic NED* is a self-contained box which performs all the functions needed by our model. It intercepts the user traffic, performs user authentication, connects to a repository for retrieving the user's security policy and applications, executes the applications locally and finally provides feedback to the user about correct processing of his network traffic. This model is suitable for small environments, with a low number of users (as the computational capacity is limited): typically a home gateway (to protect all the various devices of the family and the occasional visitors) or an edge router for a small or medium enterprise (depending on the number of simultaneously active users).

On the other hand, the *split NED* is designed to use two different logical components: the network front-end and the computational back-end. The first one intercepts the user traffic, performs user authentication, and connects to a repository for retrieving the user's security policy and applications. However, these applications are not executed locally (due to limited computational capacity or high number of simultaneous users), rather suitable computational nodes (such as those of a cloud computing system) are selected for execution of the user's applications. This model is suitable for large environments such as corporate networks (that would probably use a private cloud for the computational part) or mobile network operators (that could integrate the computational part in their NFV platform).

| *market description* | *potential users* | *expected pricing* | *potential market value* |
|---|---|---|---|
| Customers (families) interested in a parental control software | 1,383,180 | 30 €/year | 41.5 M€ |
| Customers interested in personal security software (e.g. personal firewall) | 18,281,920 | 15 €/year | 274.2 M€ |
| Corporate mobile users interested in mobile protection software | 787,913 | 50 €/year | 39.4 M€ |
| Corporate users that adopt the BYOD paradigm, interested in mobile protection software | 2,412,964 | 50 €/year | 120.6 M€ |
| Companies interested in operator-based protection software (e.g. corporate firewalls) | 52,414 | 1000 €/year | 52.4 M€ |
| *total* | | | *528.1 M€* |

*Table 1: possible yearly revenues for an Italian telecom operator from new in-network security services*

## 6. Business Benefits

The proposed architecture has several potential benefits for various categories: end-users, network operators, service providers and application developers.

### *6.1 – Benefits for stakeholders*

End-users will have the option to reduce their costs for protection by buying a single application and using it to protect all their devices, rather than buying a different application for each device. Additionally, by specifying a high-level policy (automatically translated to the configuration for a specific application), the users will be able to migrate from one application to another without the hassle of learning the different configuration options, while at the same time offering to the network operator the freedom to choose the best set of security applications for its infrastructure. Last but not least, high-level policies are also ideally suited to create standard "policy sets" that can be provided to the users by specific interest groups (such as parents' associations or user groups) thus simplifying the spread of correct and well-thought-out security configurations.

Network operators can exploit their new SDN and NFV architectures in conjunction with network offloading to reduce the time to market for new security services and to offer end-users one additional service: ubiquitous access to a NED. Of course operators providing multiple access technologies (e.g. xDSL at home and 3G/4G for mobility) in a single bundle are better positioned to offer generic access to the NED than those operators managing a single access technology; these last ones will probably need to enter into an agreement with operators handling complementary technologies.

On another hand, service providers could offer a NED independently from the network operators, by placing them strategically inside the transport networks of several operators. In this way there would be a security offer independent from the network subscription.

Finally, developers of security applications would have the benefit to have their applications available for any device rather than spending effort for porting them to the different software platforms. Hence they could better spend their time in improvements than in tedious porting activities. Moreover, the concept of generalized applications and generic policies would naturally support the creation of a marketplace for NED security applications, potentially increasing competition, lowering prices and creating a new business ecosystem as it happened in the past for mobile apps.

*6.2 – A quantitative evaluation of business opportunities for network operators*

The possibility to install security applications in the network allows telecom operators to exploit new business opportunities by selling secure ubiquitous access to their customers. For instance, a recent study [9] of the Italian market (which accounts 41.37 M Internet users and 60.9 M inhabitants at December 2012) evaluates the potential customers and the corresponding market value originated by five possible types of security services. The results (summarized in Table 1) show that those services alone have a potential market value for the telecom operators of more than 500 M€ per year, on a global market for telecom services estimated at about 60 billion €/year. The personal opinion of the authors is that these numbers can be easily reached if a strong marketing campaign will be carried out, focusing on the possible service (e.g. "Your kids will be always protected wherever they go and whatever they use to connect to the Internet") instead of the underlying technology which, for most of the people, is an unnecessary detail. In this respect, the most critical aspect from the business perspective is the availability of easy to use (yet powerful) *applications* that allow experiencing the new service.

Although this preliminary quantitative analysis does not take into account the cost of delivering the service, nor the difficulties to turn the potential value into actual revenues (i.e. convincing the potential customers to buy the new services), it shows that the possibility to deliver ubiquitous in-network security services can offer to the telecom provider interesting business opportunities.

*6.3 – Efficiency enhancements for network operators*

Security application offloading can offer relevant efficiency enhancements to network operators. These enhancements can be directly translated into several cost reductions at different points.

First of all, network operators can migrate from heterogeneous security appliances, usually requiring specific management mechanisms for each device (e.g. a Firewall appliance or Security Network monitoring console), to a unique normalized device supporting a unified remote management (the NED and its interfaces).

Another relevant cost reduction is associated to service configuration, the management of incidents to resolve and the continuous updates and security patches that any security network tool requires. These entire tasks can be managed in the NED as a network-centric model and integrated into the management flows network operators apply to provision and control their network infrastructure, opposed to manage devices at client premises, which usually implies limited remote access or physical displacement.

Finally the Split NED model will provide the very same benefits that NFV promise in general, such as the growth of on-demand availability or the extensive use of general-purpose hardware.

## 7. Market strategies

Each stakeholder has different ways to benefit from the security offloading model proposed in this paper.

Developers potentially interested in this model range from freelancers or open-source based companies up to high-level security software enterprises. The first ones could start with small security applications, based on legacy or open-source code, with revenues coming from different models, such as dual-licenses or professional services. Some could even evolve to new disruptive security services easily to deploy based on this model. Big players could focus in creating or migrating carrier-grade or high-end professional security

software oriented to Virtualised Network Function (VNF) [10] that will also take advantage of SECURED benefits with an easy adaptation to potential mass market.

Hardware vendors are an additional kind of stakeholder that could join this model. As SDN and NFV deployments increase within different businesses due to the use of cost-effective commodity servers and industry standard software, there is also a growing need for standardised and open networking hardware to support such modern converged infrastructures. These converged infrastructures would combine all the networking, compute and storage into a seamless and highly configurable entity, which would in turn support the configurability and security requirements needed by NEDs and application offload. Furthermore, there are potential opportunities for software developers that compete in hardware market with security appliances. Business case could emerge for NED implementation as home gateway or enterprise UTM (Unified Threat Management) appliance that include the NED in hardware and a marketplace with different security appliances to compete in the security market.

Service providers and mainly network operators can exploit the user-centric approach described above to create business opportunities around user security needs, in contrast with network-centric classic models that offer network security applications that oversize the needs of users offering complex security devices. In this respect, two alternative strategies can be foreseen: the on-premise and network models.

The on-premise model implies to include the monolithic NED capacity in the network device. It may be a success strategy for medium-sized service providers or network operators with an internal process for distributing their own CPEs *(*Customer Premises Equipment), allowing them to integrate the NED as part of the home gateway. For SME clients a network operator can rent or sell dedicated "NED-enabled" devices, where security services are usually provided within the security device (firewall or network device). This approach can reduce the portfolio of physical security devices, while keeping the service portfolio allowing cost reduction per device.

The second alternative, the network model, will focus the NED deployment on the network infrastructure itself. This is the natural model for those having their own network infrastructure or wherever on-premise options are not feasible (mobile network operators). We propose a possible strategy to be followed by a network operator defined in progressive steps. As a first step, start with legacy network access devices (BRAS, PGW [11], etc.) that allow redirecting traffic to NED devices to save in equipment and to avoid disruptions, while in parallel new equipment acquisitions will include the NED platform In a second step, the operator would go through a progressive migration of legacy security services, like parental control or managed security, adding new clients to the NED platform and offering the same or even improved services to stimulate wider interest and accelerate migrations. Finally the last step on this trip will bring the growth in capacity based on the NFV model, with an infrastructure based on a dedicated cloud with COTS servers and cloud management and orchestrations systems fully integrated with network operator and service provider internal processes: Technical validation, integration with OSS/BSS environments, end user public portals, etc.

## 8. Conclusions

We have described in this paper motivation, architecture and technologies to support executing security applications into a trusted and secure network node rather than at an end-user device. This research is part of the SECURED project (http://www.secured-fp7.eu), co-funded by the European Commission under the ICT theme of FP7 (grant agreement no. 611458). This project has currently designed the architecture and developed the base components, and will deliver a first prototype on October 2014, to demonstrate the

feasibility of the proposed approach, while a set of more mature components will be gradually released in the following years.

We think that providing high protection to all personal devices independent of their capabilities and network connection is an important issue, especially given the increasing number of networked personal devices used by most people. Achieving this result with an architecture which is in-line with Future Internet evolution (SDN and NFV) is surely a bonus as it enables various business opportunities.

## References


[1] Google Glass, http://www.google.com/glass/

[2] "S-L-O-W Android Updates = Refunds and Release from Contracts?", http://www.wirelessandmobilenews.com/2013/04/s-l-o-w-android-updates-investigated.html

[3] Trusted Computing Group "TPM Main Specification, Level 2, Version 1.2, Revision 116", 2011, https://www.trustedcomputinggroup.org

[4] "TrouSerS - The open-source TCG Software Stack", http://trousers.sourceforge.net/

[5] R. Sailer, X. Zhang, T. Jaeger, L. van Doorn, "Design and Implementation of a TCG-based Integrity Measurement Architecture", 13th USENIX Security Symposium, San Diego (CA, USA), 2004

[6] F. Armknecht, Y. Gasmi, A.-R. Sadeghi, P. Stewin, M. Unger, G. Ramunno, D. Vernizzi, "An Efficient Implementation of Trusted Channels Based on OpenSSL", 3rd ACM Workshop on Scalable Trusted Computing, Alexandria, (VA, USA), 2008

[7] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, J. Turner, "OpenFlow: enabling innovation in campus networks", ACM SIGCOMM Computer Communication Review, 38(2), March 2008

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

[9] F. Risso, A. Manzalini, M. Nemirovsky, "Some Controversial Opinions on Software-Defined Data-plane Services," 1st Workshop on Software Defined Networks for Future Networks and Services (SDN4FNS), Trento (Italy), November 2013

[10] ETSI GS NFV 003, "Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV", http://www.etsi.org/deliver/etsi_gs/NFV/001_099/003/01.01.01_60/gs_NFV003v010101p.pdf

[11] 3GPP TS 23.002, "Network Architecture", http://www.3gpp.org/DynaReport/23002.htm


 