An SDN/NFV-enabled Enterprise Network Architecture Offering Fine-Grained Security Policy Enforcement

Claas Lorenz
claas.lorenz@genua.de

David Hock
Johann Scherer
{hock,scherer}@infosim.net

Raphael Durner
Wolfgang Kellerer
{r.durner,wolfgang.kellerer}@tum.de

Steffen Gebert
Nicholas Gray
Thomas Zinner
Phuoc Tran-Gia
{steffen.gebert,nicholas.gray,zinner,trangia}@informatik.uni-wuerzburg.de

August 29, 2016

Abstract

In recent years, the number of attacks and threat vectors against enterprise networks have been constantly increasing in numbers and variety. Despite these attacks, the main security systems, for example network firewalls, have remained rather unchanged. In addition, new challenges arise not only to the level of provided security, but also to the scalability and manageability of the deployed countermeasures such as firewalls and intrusion detection systems. Due to the tight integration into the physical network's infrastructure, a dynamic resource allocation to adapt the security measures to the current network conditions is a difficult undertaking. This article covers different architectural design patterns for the integration of SDN/NFV-based security solutions into enterprise networks.

1 Introduction

The security system which commonly forms the first line of defense in today's enterprise networks consists of a Perimeter Gateway Firewall (PGF). Positioned at the edge of the network, the PGF inspects and filters all incoming and outgoing packets according to the configured security policy. Traffic spikes are often handled by over-provisioning, resulting in increased operating costs. Furthermore, this approach provides no protection against attacks conducted by malicious or previously compromised nodes inside the network. Therefore, additional security systems have to be installed at every security boundary, which results in high acquisition and maintenance expenses. Often, this leads to an abandonment of these systems and an implicit in-prizing into the enterprise's risk management.

As a relief, security systems based on the concepts of Software Defined Networking (SDN) and Network Function Virtualization (NFV) have been proposed to enhance overall network security while simultaneously reducing operational costs [3]. SDN separates the control from the data plane and hence allows the network operator to automatically steer individual flows via a central programmable interface [8]. This allows a fine-grained security policy enforcement and thus improves the overall network security. NFV enables the migration of typical middlebox hardware such as load balancers and firewalls into software running on virtual machines [5]. These instances can be scaled up and down, depending on the actual resource requirements without over-provisioning. Hence, combining these two technologies can provide a solid foundation for the creation of an omnipresent and scalable security solution.

Following this proposition, we analyze the transformation of enterprise networks consisting of separated cloud, network, and security components to an integrated solution enabled by SDN and NFV. To illustrate the advantages and disadvantages, we examine stateful firewalling as an example of different integration approaches. Finally, we discuss the newly introduced challenges and outline possible solutions.

2 Architectures of Traditional Enterprise Networks

We begin with a description of the status quo of traditional network architectures with special emphasis on the management systems involved as they provide the central functionality for all operational concerns. In general, management of enterprise networks includes the topics fault, configuration, accounting, performance, and security. To address these areas, network operations mainly rely on three separated columns as illustrated in Figure 1 – A Network Management System (NMS), a Cloud Management System (CMS) and a Security Management System (SMS).

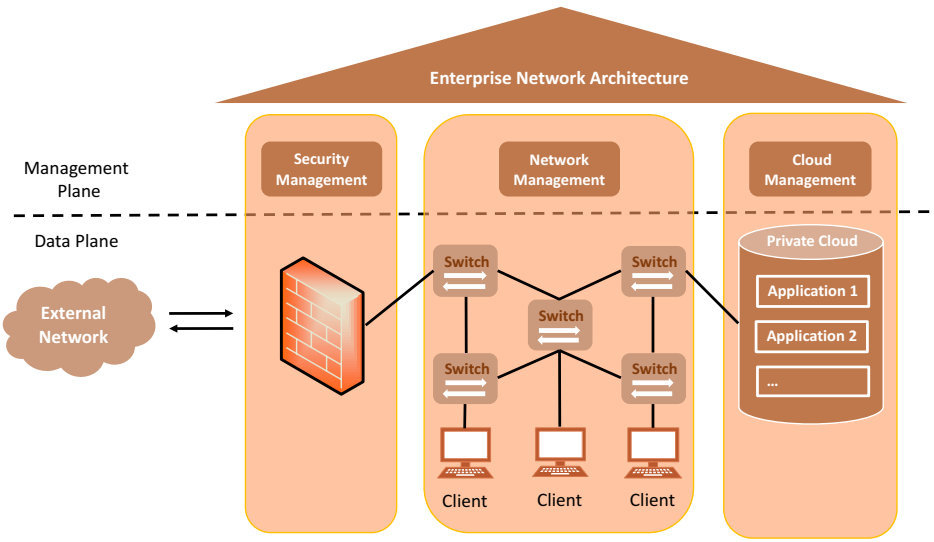

Figure 1: Traditional enterprise network architecture.

At the core, the NMS is responsible for provisioning, configuring, and monitoring the available network resources. In the provisioning stage, it provides the initial network connectivity to a newly deployed device and prepares all necessary prerequisites for the configuration stage. This can range from enabling network boot up to the installation of an entire operating system. Once this stage has

completed, the configuration stage takes over, installs the required software and applies the appropriate settings. Finally, as soon as the device is operational, the NMS transitions to the monitoring stage and triggers alarm notifications in case of detected failures.

By analogy with the NMS, the CMS plays a similar role with the distinction of having its focus set on the private cloud environment of the enterprise network. Hence, the main liabilities of the CMS reside in provisioning, configuring and monitoring of (virtual) servers and services deployed in the local cloud environment.

To address security concerns and to provide a holistic security strategy, the SMS is deployed in addition to the NMS and CMS. Its broad spectrum of tasks ranges from managing Virtual Private Networks (VPNs) to the provisioning of public key infrastructures. Having the core task of supporting the enforcement of the enterprise's security strategy, a key feature of the SMS is the ability to define and establish these policies within the network infrastructure.

As illustrated in Figure 1, the different management systems are loosely coupled. While the NMS and the CMS exchange runtime data and may utilize each other's services, the SMS remains rather isolated. This circumstance stems from the *need-to-know* principle which helps to prevent information leakage useful to an attacker. The physical deployment of security middleboxes, such as the PGF located at network borders, allows the SMS to remain independent from the other management systems. A tighter integration of the SMS with these systems, for instance by providing monitoring information, could enhance the overall operation. Similar efforts have already been conducted for CMS and NMS. In order to satisfy the demands of new services and applications, network and cloud orchestration have been integrated in the evolving service-based architecture. Therefore, it is expected that the isolated operation of all management systems will no longer be possible in the near future.

Albeit this architecture has proven to provide a functional management of the network infrastructure, it also imposes several shortcomings such as scalability issues and an increased management overhead. To mitigate these effects, we propose an enhanced enterprise network architecture based on SDN and NFV, which is described in detail within the next section.

3 Integration and Advantages of SDN/NFV-based Security Systems

Figure 2 shows an SDN/NFV-enhanced enterprise network architecture. Whereas NFV enables the operation of network nodes such as load balancers and firewalls as virtualized entities, SDN decouples the data from the control plane. Packet forwarding (data plane) is handled by SDN switches, while the SDN controller implements the control plane and decides about traffic forwarding. The SDN controller runs as a software on server hardware and provides a central interface to the network, thus offering enhanced monitoring capabilities in addition to the possibility of dynamic packet re-routing and manipulation.

Due to its central role, the SDN controller has to interact with all the other management systems. For instance, it interacts with the NMS and CMS when provisioning network connectivity and, vice versa, the NMS and CMS rely on vital network statistics conducted by the controller. In addition, the SMS needs to interact with the controller to enforce the security policy. Thus, the classical requirements of the SMS are widened to ensure the security of the virtualized entities initiated by the CMS and the SDN controller. Hence, the SMS must have control over all security related aspects of the intended operation and therefore suitable interfaces to supervise the secure execution of these processes must be provided to the SMS.

Enhanced scalability. One advantage of following this approach is the possibility to easily scale the deployed security systems according to the network load. In contrast to today's security systems, which are often implemented in costly middlebox appliances and deployed in a physical

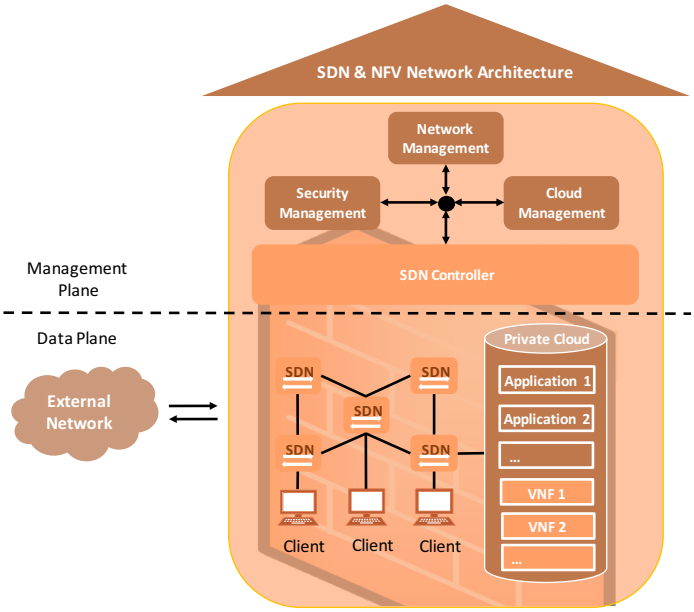

Figure 2: SDN and NFV enhanced network architecture.

stationary position, virtualized network functions run on Commercial-off-the-Shelf (COTS) servers and can be instantiated or migrated with relative ease. In particular, traffic spikes can be managed by dynamically deploying additional instances of the stressed system and, once the load has returned to normal thresholds, these instances can be discarded. Therefore, a more cost efficient resource allocation can be achieved, which in return results in lower operating and acquisition costs.

Finer granularity. In addition, enterprises often introduce a separation of duties to their divisions, which is also reflected in their security policy, such as access to human resource data is limited only to a specific group of people. In today's networks, this can be achieved on the network level through physical separation or coarse-grained logical separation using VLANs, as well through access control mechanisms on the application layer. With its fine grained flow handling, SDN offers means to dynamically define virtual networks imposed by the security policy. Each virtual network mirrors a security clearance and denies unauthenticated access. This mechanism may also be used to implement a solid policy for use cases like Bring Your Own Device (BYOD) [6], as it allows to lock a formerly unknown device into its very own virtual network with minimal access to the enterprise network. After successful authentication, its virtual network is updated to correspond to the user's security clearance and network access to further services through the means of SDN and NFV.

Flexible service chaining. Depending on the requirements imposed by the overall security policy, NFV allows to chain different security measures in a service oriented manner. For instance, parts of the traffic can be dynamically routed through a firewall, then through an Intrusion Detection System and finally through a function performing a virus scan. This flexibility allows to create a solution tailored to the needs of the enterprise, as additional services can be inserted or removed at any arbitrary position within the forwarding graph [5]. Furthermore, this decision can be made on a per-flow basis and, thus, provides advantages over statically wired classical networks.

Improved firewalling. Firewalling can be regarded as one of the most challenging aspects of the security enforcement, as it involves an active intervention into the end-to-end semantics of communications. Especially, more advanced filtering techniques like stateful firewalling typically lack

hardware support, preventing security enforcement at line rate. Yet, the use of hardware for the less advanced but common stateless firewalling in appliances is quite costly, due to the complex development cycle of specialized hardware. Today, many SDN controllers have firewalling applications included, which leverage the ability of the SDN switches to drop or forward flows. This allows for mimicking the behavior of a typical stateless firewall. In this context, the integrated forwarding tables are used as a hardware accelerator and hence establishes the first step to a cost effective and solid field firewall which relies on SDN principles and COTS hardware.

In the next section, we discuss multiple approaches based on the previously described SDN/NFV-enabled architecture to achieve a more advanced, scalable and cost effective security policy enforcement in enterprise networks featuring stateful firewalls.

4 Use Case: Stateful Firewalling

In this section, we propose three different approaches to implement stateful firewalls with SDN and NFV, as shown in Figure 3. The majority of previous work regarding novice firewall implementations either completely relies on the programmability of SDN devices, or implements the firewall functionality entirely in software. Yet, most approaches are limited to stateless firewalling, meaning that a fixed set of rules either allows a packet to pass or to get dropped.

In contrast, this work proposes a hybrid approach to implement a comprehensive, stateful firewalling concept. This includes not only that a permitted egress connection results in incoming response packets being passed, but also that the state of the connection conforms to the protocol, i.e. initiated by a TCP three-way handshake. In the following sections, we detail how challenging the tracking of the connection state in regards to the protocol's state machine can be and how this impacts the performance of different implementation approaches. Table 1 gives an overview of advantages and drawbacks of these different approaches and a comparison to classical PGF appliances. It can be seen that the hybrid approach combines the advantages of the controller-centric, SDN-based approach and the VNF-centric approach, which implements functionality in software. In contrast to a classical PGF, the proposed implementation offers better scalability and more flexibility. In general, any network function consists of parts which can be categorized as control plane functionality, for example the connection state, and other parts, which can be grouped as data plane functionality. In the case of a firewall this is the forwarding of packets. In this context, the presented approaches differ significantly in how data and control plane are constituted.

4.1 Controller-Centric Approach

The basic idea of the controller-centric approach is to use the means of SDN switches for implementing the data plane firewall functionality. In an SDN-enabled network, packets that should not be forwarded can be dropped directly inside the switches by defining flow rules matching the corresponding flows. Consequently, stateless firewalls are integrated in SDN controller software like Floodlight or FlowGuard [7]. As the actual state of the connections cannot be tracked within the switches, such control plane logic has to be implemented in the controller software. Therefore, the switches are instructed to send unknown traffic through their control channel to the controller, which also holds all security policies that should be applied. As soon as the control plane firewall function decides to forward a packet, it is sent back to the switch along with adequate information like which interface the packet should be emitted. This detour through the controller is taken until the connection is established (Path 1. in Figure 3). Afterwards, the controller installs appropriate rules on the switch that match the relevant header fields of this specific connection, i.e. the TCP five tuple. From this point on, the forwarding is handled by the switch hardware and can happen at line rate (Path 2. in

Figure 3). Hence, there is no need for further involvement of the controller. At last, the flow table entry will automatically be discarded, once the connection is inactive for a specified time.

The most significant drawback of this approach is the high latency during the connection establishment caused by the interaction between data plane and control plane through the slow control channel. Further, the load on the SDN controller is increased by taking over the data plane firewall functionality during the connection setup. Even though modern controller implementations can run as a cluster and scale out with increasing load, overloading the controller with more and more network functions has to be avoided. Another problem is the very limited space within the hardware flow tables of switches, which is described in Section 5 in detail.

Approach	Description	Pro	Contra
Controller-centric	Handshake handled by the controller.	• Follows SDN-principles • High throughput for established connections	• High Latency during connection setup • Does not scale well
VNF-centric	All traffic is diverted via firewall VNFs.	• Low latency during connection setup • Good scalability and reliability • Possibility for application level filtering	• Limited throughput per instance • Higher resource usage through multiplication of traffic
Hybrid	VNF-centric for connection setup. Controller-centric for long lasting and data intensive connections.	• Good scalability • Low latency • High throughput	• High complexity • Application layer filtering not for all connections
Classical PGF Appliance	An appliance is placed between two networks and all traffic flows through it.	• Physical placement enforces filtering of all traffic flowing between two networks • High ability for self defense possible • Application layer filtering for all connections	• Can hardly handle virtual networks due to physical placement • Very high costs per instance • Limited throughput per instance • Does not scale well

Table 1: Advantages and drawbacks of different stateful SDN/NFV firewalling approaches and classical PGF appliances.

4.2 VNF-Centric Approach

To mitigate limited scalability and high latency, the VNF approach relies on virtualized firewalls which are deployed in a cloud environment. For this, all traffic is diverted to the firewall Virtualized Network

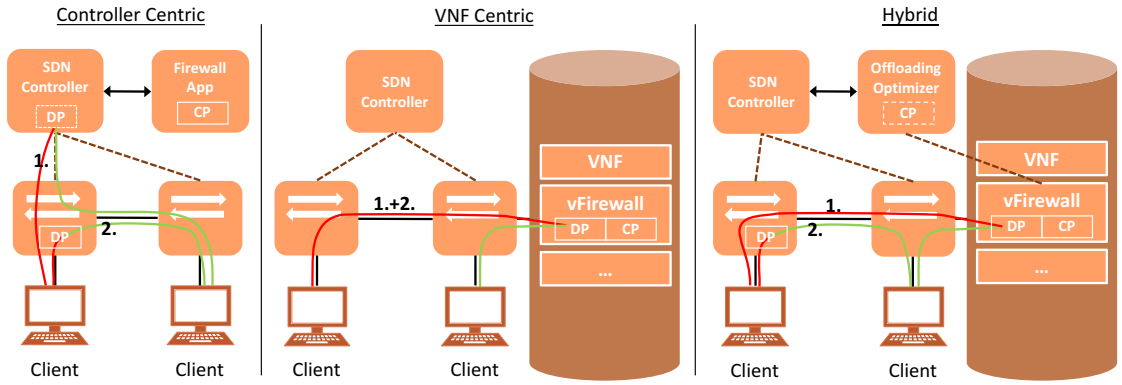

Figure 3: Three approaches implementing stateful firewalling with NFV and SDN.

Functions (VNFs), which implement tracking of the connection state as well as filtering and traffic forwarding. Therefore, the control plane (CP) and data plane (DP) of the firewall are united in one entity, which is similar to legacy firewall appliances. In consequence, all traffic (1.+2.) is routed via the VNF as shown in Figure 3.

Software-based firewalls already existed on the market before the introduction of the NFV concept and are available as commercial products as well as open source software, including Cisco's ASAv and pfSense or IPFire respectively. These products are not in line with the NFV concept, as they lack the ability to scale out and instead only operate in active/passive setups, where one instance handles all traffic. Firewall implementations following the VNF model [4] and utilizing mechanisms of cloud applications illustrate the benefits of NFV in terms of scalability and reliability compared to traditional deployment models. Further, the softwarization of network functions enables a scenario-tailored deployment of instances and function blocks on the available COTS hardware. For the discussed firewall use case, advanced filtering capabilities can be implemented and incorporated by intrusion detection software and application layer firewalls.

The downside of a VNF implementation is the limited throughput per instance, as all processing happens in software. Additionally, the virtualization overhead and resource sharing inside the physical system result in scheduling delays, which increase forwarding delays. A multitude of optimization techniques [1] are available to increase both throughput and delays of VNFs running on general purpose hardware, however without reaching the performance of appliances based on Application-Specific Integrated Circuits (ASICs). The additional detour of the traffic passing through the firewall can have an even bigger influence when all communication is routed to a VNF running multiple hops away in the data center. At this point, special emphasis has to be put on the function placement [10], in order to avoid lengthening of traffic paths which results in a service degradation. Finally, the distributed nature of multiple stateful firewall instances requires mechanisms to synchronize necessary state information, for instance all established connections.

4.3 Hybrid SDN/NFV Approach

To overcome the drawbacks of the previous approaches, we suggest a hybrid solution that introduces a strategy to offload bandwidth-intense connections to the SDN switching hardware. This results in multiple data and control plane instances: The hybrid approach uses the switch hardware as data plane, while the still existing VNF keeps both – control state and forwarding functionality.

At first, all traffic is redirected to firewall VNFs running in the local cloud infrastructure (Path1. in Figure 3). Any VNF is solely responsible for the connection establishment and thus this results in

reduced initial latency compared to the controller-centric approach. The VNF will ensure the accordance of the connection with the security rules, the protocol and optionally even with the application layer headers.

Once the connection is persistent it may be offloaded in a second step by installing forwarding rules in the switches and thus the switch is used as firewall data plane (Path 2. in Figure 3). This decision is undertaken by an optimizer which keeps track of all active connections held by the VNFs and therefore implements the control plane functionality. If an offloading decision for a certain flow is made, consecutive packets of this flow can be directly forwarded through the networking hardware, instead of detouring through the VNF. The advantage is a lower latency, a potentially higher throughput of the traffic flow, as well as a lower resource consumption by the VNFs. Good candidates for offloading are all traffic flows that are known to be long lasting and data intensive, like large file transfers. In contrast, short-lived flows, like DNS requests, might never be offloaded to the hardware.

Therefore, the hybrid approach combines good scalability with low latency for initial and consecutive packets and a high throughput. On the one hand, resources are preserved both in flow tables of switching hardware, as well as in the computing infrastructure. On the other hand, the hybrid approach is clearly more complex to implement and to monitor, as the combination of two distinct systems complicates the development and operation. This is exacerbated by the necessity to implement state synchronization mechanisms similar to the ones discussed for the VNF-based approach.

Approach	**State Tracking**	**Connection Setup**	**Filtering Decision**
Controller-centric	Controller	CP (Software)	DP (Hardware)
VNF-centric	VNF	DP (Software)	DP (Software)
Hybrid	VNF & Optimizer	DP (Software)	DP (Soft-/Hardware)

Table 2: Placement patterns for different parts of the firewalling functionality.

To summarize, Table 2 compares the three approaches regarding where the particular functionality resides. It shows that the beneficial performance behavior of the VNF-centric and the hybrid approaches stem from keeping the state tracking and the connection setup in the data plane and thus, locally. When offloading connections, the hybrid approach propagates the filtering decision to the switching hardware, which reduces the overall workload drastically. This helps to keep the amount of required VNF instances to handle the overall traffic low and therefore helps to increase resource efficiency.

5 Challenges for SDN/NFV-based Firewalling

Despite these advantages, a novel network security architecture also imposes new challenges that need to be taken into account before deployment. As no perimeter firewall is deployed in the presented approach, network security mainly relies on the SDN infrastructure. Due to its central role, the control plane is key to the security of the complete network. On the other hand, the fine granularity introduced with this concept leads to more traffic, being filtered by the firewall and thus additional load. In the following, we will discuss the implications both on security and performance and summarize potential solutions.

5.1 Control Plane Security

With the introduction of the SDN controller as a new critical component, as well as the APIs offered by switches, new threat vectors are imposed. Hence, if an attacker gains access to the controller the

entire network is compromised. Further, controller software projects are based on a large code basis and as with every feature-rich software, vulnerabilities in the control framework itself are probable. This can be mitigated by proper quality assurance mechanisms as incorporated in most larger software projects.

Vulnerabilities residing in SDN applications that influence the controller behavior and either run as controller plug-in or communicate with the SDN controller impose a threat to the network, as these APIs usually allow the controller to alter the state of the network. Hence, a proper authentication and authorization management for such SDN applications is mandatory. Current controller software lacks access restrictions allowing the limitation of access to the controller's API.

Furthermore, conventional firewall access rules which either deny or grant access for certain flows based on their packet headers are no longer sufficient. Whereas these rules are sufficient in classical networks, they can be circumvented if SDN-enabled devices dynamically rewrite packet headers to cross security boundaries. FortNox [12] is the first publication describing this problem and offers a possible solution: An SDN controller kernel validates all requests coming from SDN applications against defined security policies and keeps track of all existing rules in the SDN switches. However, this solution requires massive changes to existing control plane software frameworks and continuous work to keep track with their ongoing development.

The communication between switches and controllers opens another vector to a potential attacker. Related work [2] reports eavesdropping or denial of service attacks. As these kinds of attacks cannot be sensed at the control plane level, security improvements to the controllers fail to detect them. To mitigate this threat, a crucial security measure is the activation of encryption mechanisms for the controller switch connections including mutual authentication.

Another possible solution for detecting malicious rules is a control plane firewall placed between controller and switches. Figure 4 shows the logical placement of such a firewall.

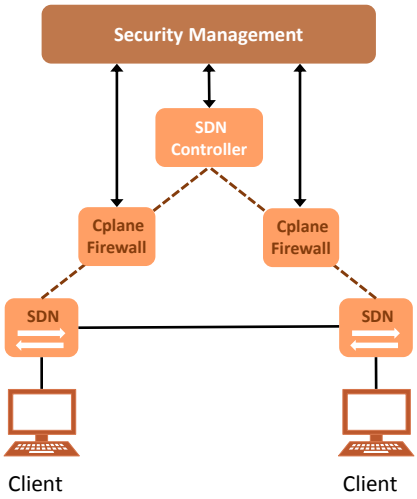

Figure 4: Control plane firewall implemented as southbound connection proxy.

This lightweight proxy instance can be positioned within the southbound communication of the controller and the switches and therefore provides a direct access to the network state as well as the forwarding hardware to the SMS. The intermediate placement requires a change of the end-to-end semantics of the controller to switch relation. The control plane firewall can be regarded as secure endpoint of the control plane connections.

The placement in the control connection enables the firewall to detect and prevent the circumven-

tion of security restrictions by rogue controllers or hostile switches. In addition, this control plane firewall cannot only protect from dangers originating from the data plane but also check the instructions that the controller forwarded initiated by calls to the controller's Northbound API. Together with the SMS as a centralized entity, a security cage is established around the SDN controller and therefore provides the necessary network-wide flow validation.

5.2 Performance Limitations

Adjacent to the described security considerations, the proposed architecture also imposes performance challenges. Briefly approached in Section 4, these challenges are described in this section in more detail. In general, SDN is assumed to improve network performance and utilization, as packet processing is done in hardware by the switches which can offer further features, like load-balancing and network virtualization based on the controller's instructions. Yet, the amount of space in the flow tables is limited and differs from model to model. This diversity in SDN hardware also causes other problems including differing forwarding delays as shown in [11]. In particular, a firewall relying on 5-tuple rules to be installed in the switches can therefore restrict the scalability of such implementations. To overcome this issue, a careful aggregation or a clever use of these resources is necessary. A related issue are delays originating from the control to data plane interaction. The authors of [9] show that the flow setup rate in an SDN environment is limited to less than one rule per ms. Especially the hybrid approach presented in Section 4 can be a relief, as only a limited number of flows is installed in the hardware forwarding table. Thus, the offloading decision is a major challenge regarding this approach. A simple solution could be to offload flows after a fixed duration or based on implicit knowledge on the applications in the enterprise environment. Nevertheless, limited resources in switches and for the VNF must be kept in mind, and therefore load balancing between VNF and switch may propose a viable solution. Up to now, no research results were presented regarding this problem.

Another challenge that can be identified is the limited support for higher layers in SDN hardware. In contrast to most network functions, like load balancers, a modern application layer firewall examines different flows very carefully and even filters content, such as HTTP traffic. As these functionalities are implemented in software, this requires large multi-purpose computing capabilities which are not provided by current network hardware (switches and routers) and thus, the implementation of an advanced stateful firewall supporting application layer filtering based solely on SDN is a difficult undertaking. On the other hand, in combination with NFV, which can be used for stateful and application layer filtering, an increase in flexibility and cost efficiency is expected.

6 Conclusion

Facing a constant increase of threat vectors, deploying and maintaining secure enterprise networks is becoming increasingly challenging and costly. As today's security mechanisms are often tightly integrated into the physical network infrastructure, the implementation of a dynamic resource allocation based on current network characteristics such as the load is a difficult undertaking. To provide a higher flexibility and to reduce operational costs, security mechanisms based on Software Defined Networking (SDN) and Network Function Virtualization (NFV) have been proposed by the research community.

In this work, we summarize how a traditional enterprise network architecture can be extended to incorporate the concepts of SDN and NFV. Furthermore, we outline the main benefits of this approach. Taking stateful firewalling as example, we illustrate three potential design patterns for the implementation, which are a Controller-centric, a VNF-centric and a hybrid approach. By discussing the pros and cons of each design pattern, we provide an overview, which can be used as guideline

for the development and possible integration of SDN and NFV appliances into current enterprise networks.

Despite the improved scalability and flexibility provided by an SDN and NFV enabled network, new challenges are imposed, which need to be taken into account. In this context, we see the additional security considerations, which are introduced by the new SDN components and their possible performance limitations as the most pressing concerns.

Yet, we are convinced that the advantages provided by SDN and NFV outweigh the disadvantages and that the additional challenges can be tackled by further research within the following years.

Acknowledgments

This work has been performed in the framework of the SarDiNe project and is partly funded by the German Federal Ministry of Education and Research (BMBF). The authors alone are responsible for the content of the paper.

References

[1] T. Barbette, C. Soldani, and L. Mathy. Fast Userspace Packet Processing. In *Proceedings of the Eleventh ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, ANCS '15, pages 5–16, Washington, DC, USA, 2015. IEEE Computer Society.

[2] K. Benton, L. J. Camp, and C. Small. OpenFlow Vulnerability Assessment Categories and Subject Descriptors. pages 151–152, 2013.

[3] M. Casado et al. Ethane: Taking Control of the Enterprise. In *ACM SIGCOMM Computer Communication Review*, volume 37, pages 1–12. ACM, 2007.

[4] J. Deng et al. VNGuard: An NFV/SDN combination framework for provisioning and managing virtual firewalls. In *2015 IEEE Conference on Network Function Virtualization and Software Defined Network (NFV-SDN)*, pages 107–114. IEEE, 2015.

[5] ETSI ISG NFV. Network Functions Virtualisation (NFV): Architectural Framework, RGS/NFV-002, V1.2.1, 2014.

[6] S. Gebert et al. Demonstrating a Personalized Secure-By-Default Bring Your Own Device Solution Based on Software Defined Networking. In *28th International Teletraffic Congress (ITC)*, Würzburg, Germany, Sept. 2016.

[7] H. Hu et al. FlowGuard: Building Robust Firewalls for Software-Defined Networks. In *Proceedings of the third workshop on Hot topics in software defined networking - HotSDN '14*, pages 97–102. ACM, 2014.

[8] M. Jarschel et al. Interfaces, Attributes, and Use Cases: A Compass for SDN. *Communications Magazine, IEEE*, 52(6):210–217, June 2014.

[9] M. Kuźniar, P. Perešíni, and D. Kostić. What You Need to Know About SDN Flow Tables. In *Passive and Active Measurement*, pages 347–359. Springer, 2015.

[10] S. Lange et al. Heuristic Approaches to the Controller Placement Problem in Large Scale SDN Networks. *IEEE Transactions on Network and Service Management*, 12(1):4–17, March 2015.

[11] A. Lazaris et al. Tango: Simplifying SDN Control with Automatic Switch Property Inference, Abstraction, and Optimization. In *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*, pages 199–212. ACM, 2014.

[12] P. Porras et al. A Security Enforcement Kernel for OpenFlow Networks. In *Proceedings of the first workshop on Hot topics in software defined networks - HotSDN '12*, pages 121–126, 2012.

Acronym Table

ASIC Application-Specific Integrated Circuit.

BYOD Bring Your Own Device.

CMS Cloud Management System.

COTS Commercial-off-the-Shelf.

CP Control Plane.

DNS Domain Name System.

DP Data Plane.

HTTP Hypertext Transfer Protocol.

IDS Intrusion Detection System.

NFV Network Function Virtualization.

NMS Network Management System.

PGF Perimeter Gateway Firewall.

SDN Software Defined Networking.

SMS Security Management System.

TCP Transmission Control Protocol.

VLAN Virtual Local Area Network.

VNF Virtualized Network Function.

VPN Virtual Private Network.

7 Biography

Claas Lorenz (claas.lorenz@genua.de) is a Security Researcher at genua GmbH in Kirchheim, Germany. He received his Master's degree in Computer Science from the University of Potsdam where he is also working on his Ph.D thesis. His research interests include firewalling, SDN/NFV and formal security verification. Since 2015 he works for the German network security specialist genua GmbH in the SarDiNe project which is funded by the German Ministry of Education and Research (BMBF).

David Hock (hock@infosim.net) is a senior consultant for research and development at Infosim GmbH & Co. KG and is coordinating Infosim's research activities. Before he was working as a research assistant at the Chair of Communication Networks at the Institute of Computer Science Würzburg where he finished his Dr. rer. nat. degree in 2014. His current main research interests are in the Unified Network and Services Management of mixed SDN/NFV, IoT, and legacy infrastructures.

Johann Scherer (scherer@infosim.net) is a senior developer at Infosim GmbH & Co. KG in Würzburg, Germany. Prior to his work at Infosim, he studied at the University of Würzburg, where he received his Master degree in 2014. His current research interests include network management and SDN/NFV integration.

Raphael Durner (r.durner@tum.de) is Ph.D student at Technical University of Munich, Germany, where he also received his Master's degree in 2014. His research interests include hybrid SDN/NFV approaches and security in SDN.

Wolfgang Kellerer [M'96–SM'11] (wolfgang.kellerer@tum.de) is a Full Professor with the Technical University of Munich, heading the Chair of Communication Networks with the Department of Electrical and Computer Engineering. Before, he was for over ten years with NTT DOCOMO's European Research Laboratories. He currently serves as an associate editor for IEEE Transactions on Network and Service Management and on the Editorial Board of the IEEE Communications Surveys and Tutorials.

Steffen Gebert (steffen.gebert@informatik.uni-wuerzburg.de) is working towards his Ph.D at the University of Würzburg Germany, where he also received his Diploma degree in 2011. His research interests include softwarized networks and agile network operations.

Nicholas Gray (nicholas.gray@informatik.uni-wuerzburg.de) is Ph.D student at the University of Würzburg, Germany, where he also completed his Master's thesis in 2015. Here, his research interests include SDN/NFV architectures and their impact on network security.

Thomas Zinner (zinner@informatik.uni-wuerzburg.de) received his Diploma and Ph.D degrees in computer science from the University of Würzburg, Germany, in 2007 and 2012, respectively. He is heading the research group on "Next Generation Networks" at the Chair of Communication Networks, University of Würzburg. His main research interests cover video streaming techniques, implementation of QoE awareness within networks, software defined networking (SDN) and network function virtualization, as well as the performance assessment of these technologies and architectures.

Phuoc Tran-Gia (trangia@informatik.uni-wuerzburg.de) is a Professor and Director of the Chair of Communication Networks, University of Würzburg, Germany. He is also a member of the Advisory Board of Infosim (Germany) specialized in IP network management products and services. He has published more than 100 research papers in major conferences and journals. He was a recipient of the Fred W. Ellersick Prize 2013 from the IEEE Communications Society.