Report from Dagstuhl Seminar 17032

Network Function Virtualization in Software Defined Infrastructures

Edited by

David Hausheer[1], Oliver Hohlfeld[2], Diego R. López[3], Bruce MacDowell Maggs[4], and Costin Raiciu[5]

1 **TU Darmstadt, DE,** hausheer@ps.tu-darmstadt.de
2 **RWTH Aachen, DE,** hohlfeld@comsys.rwth-aachen.de
3 **Telefonica I+D – Seville, ES,** diego.r.lopez@telefonica.com
4 **Duke University – Durham, US,** bmm@cs.cmu.edu
5 **University Politehnica of Bucharest, RO,** costin.raiciu@cs.pub.ro

Abstract

The softwarization of networks by introducing concepts such as Software-Defined Networking (SDN) or Network Functions Virtualization (NFV) currently massively changes network management by enabling more flexible communication networks. The main goal of this seminar was to gather researchers from academia, industry, and standardization bodies to discuss a joint perspective on research questions in the field of NFV. This report contains talk summaries, reports on the discussion groups, as well as the personal statements and main challenges contributed by the seminar participants.

Seminar January 15–18, 2017 – http://www.dagstuhl.de/17032

1998 ACM Subject Classification Network Architecture and Design, Network Operations

Keywords and phrases flexible network management, network function virtualization, software-defined networking

Digital Object Identifier 10.4230/DagRep.7.1.74

Edited in cooperation with Leonhard Nobach

1 Executive Summary

David Hausheer
Oliver Hohlfeld
Diego R. López
Bruce MacDowell Maggs
Costin Raiciu

Network management currently undergoes massive changes towards realizing more flexible management of complex networks. Recent efforts include slicing data plane resources by using network (link) virtualization and applying operating system design principles in Software Defined Networking (SDN). Driven by network operators, network management principles are currently envisioned to be even further improved by virtualizing network functions which are currently realized in dedicated hardware appliances. The resulting Network Function Virtualization (NFV) paradigm abstracts network functions from dedicated hardware to virtual machines running on commodity hardware and enables a Cloud-like network management. All of these efforts contribute to a softwarization of communication networks. This softwarization represents a significant change to network design and management by allowing the application of operating system design and software engineering principles to make network management more efficient, e.g., by enabling flexible and dynamic service provisioning.

Since the NFV efforts are currently mainly driven by carriers and standardization bodies, academic research is decoupled from the industry driven NFV attempts in redesigning network management. Due to this missing link to academic research, opportunities for groundbreaking research and broad impact in academia are currently missing out. This Dagstuhl Seminar thus gathered researchers from academia, industry, and standardization bodies to establish this missing link by fostering collaborations and joint research initiatives. Thus, a particular focus of the seminar was on identifying the diverse connections between industry driven NFV efforts and current academic networking research.

The seminar brought together 24 participants in January 2017 to discuss a potential NFV research agenda within 2.5 days. The program included different invited talks that provided an overview on selected aspects of NFV and lightning talks by each participants to provide first research questions and to sketch research directions. We summarize each talk in the remainder of this report. The main focus of the seminar was then the in-depth discussion the research areas identified in the lightning talks in several breakout sessions, which we also summarize. We closed the seminar by collecting and discussing several opinions from each participant: (i) lessons learned and surprises on NFV during the seminar and (ii) open research questions. We further collected controversial statements on NFV research and asked the seminar participants on whether they agree or disagree to each presented statement. We summarize the outcomes at the end of the report.

2 Table of Contents

3 Madness Talks

3.1 Who Am I? What Am I doing here?

Diego R. López (Telefonica I+D – Seville, ES)

Challenges:
- Performance
- Security
- Design Patterns
- Deployment Patterns
- A software network discipline
- Evolve management
- Evolve architectures
- Evolve business models
- Evolve industry culture

3.2 Personal Introduction

Bruce MacDowell Maggs (Duke University – Durham, US)

- Largest DDOS Attacks by Year
- Leverage Compromised Home Cable Modems/Routers
- Account Takeover Campaign Attack Architecture
- Attacking IP Persistence: Finance Customer

3.3 Personal Introduction

Dirk Kutscher (Huawei Technologies – München, DE)

- Performance in the presence of heterogeneity and dynamic network conditions
- Rethinking collaboration of apps, transport and forwarding
- New forwarding abstractions and SDN control for that
- Enabling dynamic computation in the network securely
- Use cases: IoT, blending VoD and live streaming

3.4 Personal Introduction

Felipe Huici (NEC Laboratories Europe – Heidelberg, DE)

Despite years of research and a whole lot of noise made about the great benefits of NFV, especially in terms of lower OPEX/CAPEX, how to provide high utilization and performance isolation on x86 servers for production-level deployments have proven so tricky that most operators resort to throwing money at the problem: each NFV instance runs in its own core, ensuring reliable performance but woeful utilization. In this brief talk I make the case for applying machine learning techniques, and in particular reinforcement learning, to develop novel, more efficient NFV resource allocators.

3.5 Personal Introduction

Gabriela Gheorghe (PwC – Luxembourg, LU)

- Cloud & secure storage
- SDN security & network troubleshooting
- Policy-based security monitoring at network level

3.6 Performance Assessment of NFV

Thomas Zinner (Universität Würzburg, DE)

The variety of different virtualized network functions as well as the availability of diverse deployment options requires new means for their performance assessment. This includes appropriate benchmarking methods and tools, but also well-suited performance models based on methods like discrete-time analysis, queuing networks or the network calculus. The resulting models and corresponding key performance indicators will then enable an optimized placement of virtualized network functions.

3.7 NFV: Where are we?

James Kempf (Ericsson – Santa Clara, US)

The Gartner hype curve is a well-known way of classifying where a technology is in terms of deployment. Right now, NFV seems to be in the “Valley of Disillusionment”. This is a period in which the initially promised benefits of the technology fail to pan out and early adopters are left trying to figure out what benefit, if any, the technology has. With respect to NFV, it has promised that it would reduce operator costs and provide much faster development of new services. In 2016, the cost per bit for operators to deliver data increased above the revenue per bit (according to analyst Tom Nolan). Can NFV help solve this problem? NFV deployments have so far not shown promise in this regard. As for services, there have been a few (SD-WAN for example), and this area seems more promising. Over the next couple of years, we will see if NFV reaches the "Slope of Enlightenment" and the "Plateau of Everyday Use", and becomes a widely deployed and useful technology for operators and enterprises.

3.8 Checking Policy Compliance

Aaron Gember-Jacobson (Colgate University – Hamilton, US)

Joint work of Aaron Gember-Jacobson, Raajay Viswanathan, Aditya Akella, Ratul Mahajan, Chaithan Prakash, Robert Grandl, Junaid Khalid, Sourav Das
Main reference A. Gember-Jacobson, R. Viswanathan, A. Akella, R. Mahajan, "Fast Control Plane Analysis Using an Abstract Representation", in Proc. of the 2016 ACM SIGCOMM Conf., pp. 300–313, ACM, 2016.
URL http://dx.doi.org/10.1145/2934872.2934876

This talk highlights two major pieces of work I have done to improve the correctness of networks: OpenNF and ARC. The former is a framework for quickly and safely transferring NF state amidst traffic redistribution. The latter is an efficient control plane verifier that uses graphs to model the network's behavior and computes properties of these graphs to check that a policy is satisfied under arbitrary failures. An interesting area of future research is verifying that networks conform to more complex policies involving stateful NFs.

3.9 Challenges in Network Functions Virtualization

David Hausheer (TU Darmstadt, DE)

Our existing work on SDN/NFV covers Software-defined Multicast (SDM), NFVI benchmarking, and application scenarios of bare metal switches. Moreover, resource models for SDN/NFV data planes and seamless elasticity in hardware-accelerated NFV, as well as VNF state migration are of interest. SDN support of CDN networks and SDN/NFV machine learning have been partially tackled as well. Novel challenges include the discussion of relevant killer applications, functions, and use cases for NFV that show the greatest benefits. Also, what's going to be the long-term impact on hardware vendors, network providers, and end-users? Moreover, what are the biggest barriers for research? Finally, is NFV just another hype? And whats going to be next after NFV/SDN?

References

1 Julius Rückert, Jeremias Blendin, Rhaban Hark, David Hausheer: Flexible, Efficient, and Scalable Software-Defined Over-the-Top Multicast for ISP Environments with DynSDM. TNSM, 2016

2 Julius Rückert, Jeremias Blendin, David Hausheer: Software-Defined Multicast for Over-the-Top and Overlay-based Live Streaming in ISP Networks. JNSM Special Issue on Management of Software Defined Networks, July 2014.
3 Jeremias Blendin, Julius Rückert, Sascha Bleidner, David Hausheer: Taking the Sting out of Flow Update Peaks in Software-Defined Service Chaining. 2nd International Workshop on Management of SDN and NFV Systems (ManSDN/NFV 2015), November 2015.
4 Julius Rückert, Jeremias Blendin, Rhaban Hark, David Hausheer: DYNSDM: Dynamic and Flexible Software-Defined Multicast for ISP Environments. 11th International Conference on Network and Service Management (CNSM 2015), November 2015.
5 Jeremias Blendin, Julius Rückert, Tobias Volk, David Hausheer: Adaptive Software Defined Multicast. 1st IEEE Conference on Network Softwarization (NetSoft 2015), April 2015.
6 Leonhard Nobach, David Hausheer: Open Elastic Provisioning of Hardware Acceleration in NFV Environments. Workshop on Software-Defined Networking and Network Function Virtualization for Flexible Network Management (SDNFlex 2015), March 2015.
7 Jeremias Blendin, Julius Rückert, Nicolai Leymann, Georg Schyguda, David Hausheer: Position Paper: Software-Defined Network Service Chaining. EWSDN 2014, September 2014.
8 Julius Rückert, Roberto Bifulco, Muhammad Rizwan-Ul-Haq, Hans-Joerg Kolbe, David Hausheer: Flexible Traffic Management in Broadband Access Networks using Software Defined Networking. IEEE/IFIP Network Operations and Management Symposium (NOMS 2014), May 2014.
9 Leonhard Nobach, Oliver Hohlfeld, David Hausheer: New Kid on the Block: Network Functions Virtualization: From Big Boxes to Carrier Clouds. Computer Communication Review. July 2016.
10 Leonhard Nobach, Benedikt Rudolph, David Hausheer: Benefits of Conditional FPGA Provisioning for Virtualized Network Functions, SDNFlex Workshop, March 2017.
11 Jeremias Blendin, Yuriy Babenko, Dennis Kusidlo, Georg Schyguda, David Hausheer: Position Paper: Towards a Structured Approach to Developing Benchmarks for Virtual Network Functions. EWSDN, October 2016.
12 Leonhard Nobach, Ivica Rimac, Volker Hilt, David Hausheer: SliM: Enabling Efficient, Seamless NFV State Migration. IEEE ICNP, November 2016.

3.10 Road towards lightweight virtual network functions

Oliver Hohlfeld (RWTH Aachen, DE)

Main reference F. Schmidt, O. Hohlfeld, R. Glebke, K. Wehrle, "Santa: Faster Packet Delivery for Commonly Wished Replies", Computer Communication Review – SIGCOMM'15, 45(5):597–598, ACM, 2015.
URL http://dx.doi.org/10.1145/2785956.2790014

Current research efforts concern the creation, execution, or the placement of network function – mainly encapsulated in VMs or containers. While these virtualization efforts enable elastic scaling and cost efficient provisioning of virtual network functions, some functions require smaller and more lightweight execution environments. In this talk, I argue that there exists a case for realizing lightweight on-path network functions that do not need fully-fledged VM-based execution environments but rather run on network elements (e.g., switches). We have created a first prototype of an application agnostic execution environment to implement intelligent on-path functions in the network. By using this prototype, we show that server application performance can be optimized by utilizing these accelerated functions executed in the network core.

References

1 Florian Schmidt, Oliver Hohlfeld, René Glebke, Klaus Wehrle: Santa: Faster Packet Delivery for Commonly Wished Replies. Computer Communication Review 45(5):597–598, 2015.

2 Amir Mehmood, Oliver Hohlfeld, Dan Levin, Andreas Wundsam, Florin Ciucu, Fabian Schneider, Anja Feldmann, Ralf-Peter Braun: The RouterLab – Emulating Internet Characteristics in a Room. 11th ITG Conference on Photonic Networks, 2010.

3 Yvonne Coady, Oliver Hohlfeld, James Kempf, Rick McGeer, Stefan Schmid: Distributed cloud computing: Applications, status quo, and challenges. Computer Communication Review 45(2):38–43, 2015.

4 Marc Werner, Johannes Schwandke, Matthias Hollick, Oliver Hohlfeld Torsten Zimmermann, Klaus Wehrle STEAN: A storage and transformation engine for advanced networking context. IFIP Networking Conference, 2016

3.11 NFV for Industrial Control Networks

Jan Rüth (RWTH Aachen, DE)

The rising demand for automation and flexible manufacturing challenges the way today's Network Control Systems are structured. Current control architectures are composed of expensive and proprietary controllers that have been tailored for a specific use case. It is at the moment impossible to flexibly change the control algorithms and configurations. Therefore, the idea of a "controller in the cloud" is promising to give the flexibility and scaling properties that cloud computing has brought to networking. However, a controller in a cloud environment hosted over the Internet cannot meet the demands of current control algorithms, namely low latency and low jitter. We therefore propose to enable flexible lightweight-VNFs that can be computed on path within the switches of an industry control network that are under orchestration of the controller in a cloud environment. This idea challenges how packet processing within the switches can be flexibilized and how control algorithms may be expressed in this domain while maintaining high performance.

3.12 Seamless Elasticity in NFV Infrastructures with Heterogeneous Hardware Requirements

Leonhard Nobach (TU Darmstadt, DE)

Network Functions Virtualization (NFV) applies the principles known from *cloud computing* to network functions traditionally running on dedicated, inflexible and expensive hardware. In our work, we delve into two major topics:

First, we argue that the increased elasticity and fast provisioning properties of NFV come to the cost of performance issues on a single VNF instance, as these commonly operate on sequentially-working processors, unlike ASIC-driven high-performance appliances. However, the inclusion of hardware acceleration (HWA) into NFV infrastructures carries the danger of bringing back rigidness and inflexibility. We therefore propose a framework for elastic provisioning of *reconfigurable* hardware. VNF instances can dynamically claim and release FPGA (or NPU) resources over the network from a pool in the datacenter, if their performance requirements increase.

Secondly, the arising advantages for elasticity and dynamicity require state transfer mechanisms (for relocation or scale-in/out, x86 or HWA) that are *seamless* – the end user does not perceive any service disruption. Existing work does not consider the high costs of seamless state transfer traffic as a critical resource, however, links of network functions tend to be higher utilized than traditional server-cloud traffic, and state transfer attempts from or to *edge* or even *fog* instances aggravate this issue. We introduce the concept of *statelets* and our *SliM* state transfer mechanism, which can increase the possible link utilization up to a factor of three compared to previous approaches.

References

1 Leonhard Nobach, David Hausheer: Open Elastic Provisioning of Hardware Acceleration in NFV Environments. In: Workshop on Software-Defined Networking and Network Function Virtualization for Flexible Network Management (SDNFlex 2015), March 2015.

2 Leonhard Nobach, Ivica Rimac, Volker Hilt, David Hausheer: SliM: Enabling Efficient, Seamless NFV State Migration. In: IEEE International Conference on Network Protocols (ICNP), November 2016.

3.13 NFV Performance Profiling and Optimization

Andreas Kassler (Karlstad University, SE)

As network functions will be virtualized and may run on any given hardware the NFVI operator owns, we need proper methods that allow us to predict the NFV performance for a given infrastructure configuration. This requires a flexible profiling framework which allows to benchmark a given virtualized NF and configure the infrastructure in a flexible way. Having available KPIs from both inside virtual NF and the hypervisors, we would be able to model and predict the VNF performance. For optimization of VNF placement, we typically do not know precisely the input to the optimization problem. For example, we cannot accurately predict the amount of resources a given VNF needs to perform its functions. From optimization theory it is a well-known fact that once parameter are allowed to deviate from the nominal values, the optimal solution may become highly infeasible one. Consequently, we need to develop fast online solution algorithm that are able to cope with the uncertainty of information in a robust way.

3.14 Bring those network functions back to the plumbing!

Fernando M. V. Ramos (University of Lisbon, PT)

With Network Function Virtualization (NFV) the network functionality of dedicated hardware middleboxes is replaced by software running in VMs in commodity servers. In this talk we propose to bring (part of) the network functions back to the network (i.e., run the functions in network switches). This is made possible today by the recent advances in network data plane programmability (enabled by switch chip architectures such as RMT-PISA, for instance). Recently, high-performance switches (e.g., Barefoot Tofino) that can be programmed using high-level languages (e.g., P4) have emerged. The strengths of the proposed solution is mainly a gain in performance, in both lower latency (the NFs can be exactly where they're needed) and higher aggregate throughput (as NFs run in switches). The main challenges are the limited memory and computation resources of network switches. Interesting research questions in this space include how much of the network function can be offloaded to programmable hardware, the implications of the additional network programmability in terms of security, and how networks composed of stateful boxes that lead to "mutable data paths" (in which the handling of a packet depends not just on immutable forwarding state, but also on state that changes at packet time scales) can be verified, tested and debugged.

3.15 Network Function Virtualization in Software Defined Infrastructures

Michael Scharf (NOKIA – Stuttgart, DE)

Software Defined Networking (SDN) is an essential part of dynamic enterprise services, which include, amongst others, Network Functions Virtualization (NFV). Existing Carrier-SDN solutions provide network function management, analytics and assurance, as well as dynamic management and control. These Carrier-SDN solutions are already deployed in operational networks. Still, open research issues remain. Examples include the future design of data modeling languages (e.g., YANG, TOSCA) or verification of the interoperability in heterogeneous environments.

3.16 On Research Challenges of NFV

Christian Esteve Rothenberg (State University of Campinas, BR)

Critical issues to realize NFV in practice include proper performance evaluation methodologies towards predictable behavior and in support of optimized VNF placement. An open-source, collaborative approach is desirable to provide rich, reproducible platforms for VNF testing and benchmarking. Those issues are on our road-map and I hope the community will be able to effectively solve them.

3.17 NFV at the Edge

Tim Wood (George Washington University – Washington, DC, US)

NFV offers the opportunity to run high-performance network services in a flexible way. Edge clouds may become a new place to provide such services, potentially with very low latency access for users. The types of "users" connecting to such services also may change over time – not just users on phones or laptops, but autonomous vehicles, robots, smart city infrastructure, etc. These new types of users may require new types of NFV services, and may cause us to rethink the line between a network function and an application.

3.18 Utilizing Hardware Capabilities for Efficient NF Deployment

Fabian Schneider (NEC Laboratories Europe – Heidelberg, DE)

Current trends in SDN research, such as BFBA or P4, lead to extended possibilities to program network elements and NICs. Furthermore, hardware components such as GPUs or encryption cards allow for extended options to instantiate (virtual) network functions. The challenges stemming from this include, but are not limited to, (a) decision algorithms on how to split/delegate/offload sub-functions to hardware or in-network (b) programming abstraction that allows for such a functional split and (c) dynamic and online re-composition of NFs.

3.19 Dynamic Placement of Network Functions for Mitigating Internet-Scale Attacks

Oliver Michel (University of Colorado – Boulder, US)

Distributed Denial of Service (DDoS) attacks and other large-scale network attacks are common in today's Internet. They cost significant amounts of money each day by making information and services unavailable or forcing businesses to scale up their resources in the network or in the cloud while under attack. Over the past years however, mitigation techniques also become more advanced. Advanced Intrusion Detection and Prevention Systems (IDS/IPS) often are ASIC-based and can inspect large volumes of traffic at high rates. They use advanced application-layer analysis techniques, and leverage learning-based approaches to detect anomalies. We argue that wide-area SDN and NFV technologies can complement such advanced detection and filtering techniques. In particular, network functions can be placed on-demand during an attack to inspect and filter traffic at multiple locations throughout the network. These locations can be dynamically chosen to be closer to the attacker sources as opposed to placing all filtering at the victim's network gateway where the attack traffic volume can already be intractable to control. As we imagine such an architecture also in an inter-domain (AS) setting, several questions regarding a centralized trusted party and business models for cooperative, on-demand DDoS defense arise.

3.20 User-Centric NFV

Theophilus Benson (Duke University – Durham, US)

NFV has emerged as a panacea for a multitude of problems. Surprisingly, many solutions aim to provide homogeneous substrates of infrastructures. In this talk, I will argue for specializing the configuration and infrastructure to optimize the performance delivered to an end user. Although this talk focuses on performance, the broad concept of specialization can be used to improve other properties: security and reliability.

In this talk, I will present several configuration examples that impact the end user's performance. Motivated by these benefits, I will present a framework for optimizing and specializing NFV infrastructures: I will discuss inputs and constraints to the optimization problem. I will conclude by asking others to think of the potential of heterogeneity on emerging use-cases.

3.21 Flexibility as a Design Guideline for NFV and SDN

Wolfgang Kellerer (TU München, DE)

The following questions are still challenging in NFV research: How to cope with the emerging network dynamics? How to design a network for flexibility? How to migrate (network) functions? What is the role of virtualization and SDN? How can we use "flexibility" as a measure to compare different system designs?

4 Lightning Talks

4.1 An overview of network verification

Costin Raiciu (University Politehnica of Bucharest, RO) and Aaron Gember-Jacobson (Colgate University)

Networks are a critical part of our society, but managing them is tedious and error-prone. Configuration errors are common and bring major disruptions. NFV and SDN will further amplify the difficulty of ensuring networks are functioning correctly, i.e. they obey their operator policy.

The talk will cover efforts to verify formally that a network obeys the policy of its operator and has two parts focusing on the control plane verification and data plane verification. Data plane verification means taking a snapshot of the data plane state, converting it into a verifiable model; next the policy is checked against the model.

The talk will briefly cover HSA and NOD the best-known data plane verification tools, and details the design and implementation of Synet, our symbolic execution tool optimized for networks, and the associated SEFL language.

Control plane verification models the control plane, rather than the data plane, allowing policies to be verified across arbitrary failure scenarios. The talk will briefly cover Batfish, a recent control plane verifier, as well as discuss ARC, a control plane verifier that speeds up network verification by casting verification as a graph analysis problem. One limitation of control plane verifiers is their inability to capture subtle ties in the implementation of routing protocols on different vendors' devices.

4.2 Network Programmability: A Primer of Routing Synthesis

Laurent Vanbever (ETH Zürich, CH)

Between 50% and 80% of the network downtime are due to human, not equipment failing. Most of these are due to configuration mistakes, i.e. due to humans. In this talk, I present new directions in network management, one area akin to declarative programming. Specifically, I describe two complementary ways to declare and provision forwarding state, network-wide, of legacy equipment. One of the body of work is fibbing, a way to generate inputs to a distributed algorithm (Dijkstra) is that it computes what the operator wishes. In the second body of work, Synet, I explain how one can actually generate the distributed algorithm itself.

4.3 The NFV Standardization Concoction

Diego R. López (Telefonica I+D – Seville, ES)

The goal of the talk was to provide an introduction to the most relevant standardization efforts in the NFV arena, and to explore how they were related to current research challenges, and how the technology itself was shaping a new way to produce such standards.

We started with a brief introduction to the research challenges in NFV, from the perspective of a network service provider that has participated in the development of the NFV concepts since their inception. A list of technology and business challenges were introduced and briefly discussed.

The core of the talk was focused on the two main standards-focused groups currently involved in NFV matters: the ETSI NFV ISG and IRTF's NFVRG. The evolution, structure, and current plans of ETSI NFV where introduced, discussing the three types of deliverables the community aims at: normative and informative (as in many other standards bodies), and demonstrative (in the form of PoCs and early interoperability assessments)

The origin and goals of the NFVRG were introduced afterwards, detailing its research agenda, as an initial input to the discussion planned for the seminar. Later, the activities connected with NFV in other standards bodies were presented, and finally a discussion of relevant open-source projects and their influence in the standardization process was presented.

The talk concluded with a summary of how NFV is shaping standardization while it is standardized itself, and a call to the participants to bring their results to these ongoing efforts.

4.4 The Functions Placement Problem

Wolfgang Kellerer (TU München, DE)

Joint work of Wolfgang Kellerer, Arsany Basta, Andreas Blenk, Marco Hoffmann, Hajo Morper, Klaus Hoffmann

Network Function Virtualization (NFV) together with Software Defined Networking (SDN) opens up new challenges for the composition, placement and migration of network function in an operator's network. We refer to this class of problems as the Function Placement Problem (FPP) inspired from the Controller Placement Problem that has been introduced for SDN controllers by Heller in HOTSDN 2012. However, we see new types of challenges arising with the FPP. First, as part of an optimal placement an optimal function (de-)composition and chaining has to be considered. SDN and NFV offer complementing concepts here where network functions can be moved completely (based on NFV) or partially (based on the SDN control/data plane split) into a data center [1]. Second, we address dynamic placement and migration as a further important network design aspect [2]. Finally, we show how to analyze the flexibility a chosen network design to understand its overall benefits [3].

References

1. A. Basta, W. Kellerer, M. Hoffmann, H. Morper, K. Hoffmann. *Applying NFV and SDN to LTE Mobile Core Gateways; The Functions Placement Problem*. AllThingsCellular14, Workshop ACM SICGOMM, Chicago, IL, USA, August 2014.
2. A. Basta, A. Blenk, M. Hoffmann, H. Morper, K. Hoffmann, W. Kellerer. *SDN and NFV Dynamic Operation of LTE EPC Gateways for Time-varying Traffic Patterns*. 6th International Conference on Mobile Networks and Management (MONAMI), Würzburg, Germany, September 2014.
3. W. Kellerer, A. Basta, A. Blenk. *Using a Flexibility Measure for Network Design Space Analysis of SDN and NFV*. Software-Driven Flexible and Agile Networking (SWFAN), IEEE INFOCOM Workshop, San Francisco, USA, April 2016.
4. W. Kellerer, A. Basta, A. Blenk. *Flexibility of Networks: a new measure for network design space analysis?* arXiv report, December 2015[1].

4.5 Getting out of the NFV Deployment Quagmire

Felipe Huici (NEC Laboratories Europe – Heidelberg, DE)

Containers are in great demand because they are very lightweight when compared to virtual machines: both boot times and memory usage are significantly smaller than traditional VMs, and this allows massive consolidation of workloads on the same hardware. On the downside, containers have fundamentally weaker isolation properties than VMs.

In this talk, we examine whether there is indeed a strict tradeoff between isolation (VMs) and efficiency (containers). By redesigning the control plane of Xen and using small, optimized unikernel based virtual machines we show that it is possible to achieve VM boot times on the order of milliseconds while packing thousands of VMs on modest hardware.

---

[1] http://www.lkn.ei.tum.de/forschung/publikationen/dateien/Kellerer2015FlexibilityofNetworks:a.pdf

5 Working Groups

5.1 NFV Security, Validation, and Verification

Aaron Gember-Jacobson (Colgate University – Hamilton, US)

Joint work of Aaron Gember-Jacobson, Gabriela Gheorghe, James Kempf, Diego R.López, Bruce MacDowell Maggs, Costin Raiciu

- Service level agreements drive operator's efforts to verify/attest correctness, performance, etc.
- How does verification change as you go from hardware to software?
  - NFV is less trustworthy:
    - With hardware boxes, the development and deployment cycle is long: The vendor takes time to develop an NF and the operator takes time to certify the NF.
    - NFV tends to follow a DevOps model: less testing → higher likelihood of bugs → verification becomes more important. To benefit from NFV, you need to change the corporate culture within the ISP regarding the need for extensive certification and having individual boxes fulfill 99,999% (5-nines) uptime.
  - Verification of chaining:
    - With hardware, you simply trace the cables.
    - With software, it is much more difficult – you now have the chance to deploy very complex chains of small VNFs that were not deployable before, which makes the big picture more complex.
- How do we go about doing verification?
  - Lots of languages/tools for expressing/checking policies – not clear these are used.
  - P4 is valuable as a modeling language.
  - Verification could be made simpler by leveraging flexibility offered by NFV, for example...
    - . . . always put NFs in a particular order that is easier to verify
    - . . . break NFs into micro-services – introduces interoperability challenges; could also make verification harder, as noted above.
  - If you optimize for performance, then verification is hard, and vice versa.
- Operators want to be able to provide micro-containerized services that can be flexibly moved around.
  - Infrastructure needs to know if one can trust the hardware and software needs to know if one can trust the infrastructure.
  - Who vouches for the authenticity of an NF? How do you do this in real time? – we know how to do attestation, but doing it at scale is hard.
  - How do we apply attestation mechanisms for hypervisors to SDN controllers and switches?
  - How do you protect private keys stored on NFs that split TCP connections? – Can't put keys on every edge NF.

5.2 Placement

Wolfgang Kellerer (TU München, DE)

Joint work of Andreas Kassler, Wolfgang Kellerer, Diego Lopez, Michael Scharf, Fabian Schneider, Andreas Wundsam, Thomas Zinner

- Scope: Offline vs. online; the focus is online!
  1. Ways to decide to program and deploy depending on hardware capabilities, for example DPDK or P4. → Hardware offloading/transition? Heterogeneous hardware (adaptive hardware)?
     - Decomposition of VNFs:
       * Manual/Automatic,
       * Optimization,
       * Structural – designing the system,
       * Dynamic – reacting to environment/traffic.
  2. Machine Learning (control loop)
     - Actions:
       * Resource allocation, e.g. VM migration,
       * scheduling; single server (early papers exist).
     - Metrics/objective functions (different ones).
       * What metrics? – traditional (hardware), software, service level (KPI).
     - Monitoring / telemetry
  3. Heuristics for VNF placements (e.g., co-location)
     - Online and offline placement is mostly a planning process.
     - There could be placement (optimization) or assignment problems, or also other interconnected problems (e.g. from switch to controller).
  4. Dynamics – Changing the system
     - What is the cost of migration (resource, time)?
     - Verification of continuity of operation
     - Debugging / root cause analysis
     - Robustness, e.g. during state migration, uncertainty of prediction, impact?
- Open Questions / Discussion
  - Machine Learning – how good is it?
    * Where does the data/training come from? The assumption is that data is available.
    * Machine Learning <-> Heuristics / Optimizers
    * Both may may be combined, e.g. machine for pre-filtering and then heuristics or optimization of a simpler problem.
  - Reaction to unknown situations (e.g., attacks!)
  - Models and abstractions for performance
    * How realistic is a "world model"?
      · Where to place, when to place?
      · What are the parameters that we need to optimize for?
      · How is the real hardware working?
    * SLA metrics for NFV (under standardization)? KPI? Abstractions?
      · *Remark from Diego Lopez:* There is some ongoing work in ETSI on performance characterization and maybe some early work on KPI. There is nothing on SLA so far.
  - Conflicting goals → how to sort this out?

    - User: Best performance (SLA-defined)
    - Operator: "Run infrastructure hot"
- Discussion:
  - *James Kempf:* What algorithms would you use? Traditional optimization algorithms are not designed for online use.
  - *Diego Lopez:* Call to practitioners for data-driven management [seminar conclusion].
  - *Thomas Zinner:* Machine learning is not good for systems under change since the algorithms are not trained for the new conditions.
  - *Andreas Kassler:* How would one compare different heuristics? There is no standard set of agreed-upon topologies, workload, ... for the evaluation. We need something like a standardized VNF evaluation database.
  - Suggestion: This may be a topic for another Dagstuhl seminar or a project proposal.
  - *Diego Lopez:* To have predictable performance, you need plenty of knowledge about the underlying hardware/platform (ongoing work with Intel). Open-source MANO: osm.etsi.org.
  - Security properties of the network need to be reflected, not only topology.
  - *Tim Wood:* What is different about NFV placement as compared to other placement problems?
  - Network topology has more impact on how things are run.
  - NFV, by definition, has high I/O.
  - In the end it is an optimization problem, just more complex.

5.3 NFV Economics

Leonhard Nobach (TU Darmstadt, DE)

Joint work of Tim Wood, David Hausheer, Jan Rüth, Leonhard Nobach, Oliver Hohlfeld, Theophilus Benson

- ISPs provide services to over-the-top (OTT) providers to deploy services.
  - Why would they do that? Money, Network offloading.
  - Follow-me-Cloud: Which applications exist? High computing power required with low latency.
  - However, much latency added at the edge (e.g. forward error correction).
  - Privacy issues
    - Locality → local laws
    - Power over infrastructure
  - Do we need more than sandboxing via SDN?
  - Are there mechanisms to guarantee that the provider cannot e.g. snoop the memory of my VNF?
  - Asymmetric bandwidth: The reason is DDoS. Edge clouds would make it worse.
  - DDoS Detection would be a good use case: An NF may search for signs for DDoS.
  - Suggestion: An NF which returns aggregates of customer data only (data protection).
  - New economics:
    - The ISP provides feature for OTTs to install VNFs, It then charges the customer.
    - Customer requests OTT service, OTT installs VNF to deliver the better service to customer.
    - Aggregation of multiple customers possible.

    - Small and medium enterprises.
  - What is the cost/revenue share between the ISP and OTT provider?
  - Regulation: Is this net neutrality?
    - If every OTT can deploy VNFs: No need for regulation.
  - Use case for low latency: Tactile gaming.
- Discussion:
  - Objections from *Diego Lopez*
  - Idea: Tactile e-commerce (James)
    - *Bruce Maggs:* E-Commerce is the one place today where people pay a lot of money for low latency. How can you use NFV to realize that?
  - *Dirk Kutscher:* There are only two use cases for edge computing.
    - Lots of data is generated locally.
    - Interactive VR
    - For all other use cases, one is much better off reducing the overall latency.
  - *Tim Wood:* The pricing of these resources is going to be complicated
  - *Bruce Maggs:* This sort of net neutrality doesn't exist yet. Operators currently charge whatever they want
  - *Michael Scharf:* For a VNF to be useful you need 1) a VNF and 2) connectivity. This can provide some interesting economics related to location: 1) infrastructure is idle (and thus cheap) but connectivity is overloaded. 2) infrastructure is hot, but connectivity has idle resources.
  - *Michael Scharf:* The other interesting implication is the last-mile connectivity is typically the expensive part.

5.4 NFV Performance

Theophilus Benson (Duke University)

Joint work of Theophilus Benson, Christian Esteve Rothenberg, David Hausheer, Oliver Hohlfeld, Leonhard Nobach

- Fundamental questions:
  - What is a VNF?
  - How can we systematically model and measure the performance characteristics of a VNF?
  - How can we make (or develop) techniques and frameworks that enable a software (VNF) to provide similar performance guarantees as hardware-based functions?
- Understanding and framing the performance problem:
  - Types of metrics: App-specific (SIP rps, HTTP rps), user-specific (page load time), or general VNF metrics.
  - General VNF Metrics: first class (latency/throughput), second class (migration/boot-up time).
- Use cases for VNF:
  - What is a VNF: a middlebox running in a VM or a middlebox rewritten for "the cloud"?
  - Control functions – DNS, DHCP, BGP aggregation.
  - Data functions – focus on network functions that cannot be supported in ASICs/P4.

    - A stateful load balancer which requires HTTP header information,
    - A dynamic deep packet inspector (DPI): flexible regular expression or signature matching not supported,
    - Software-defined WAN: WAN optimization not supported in ASIC,
    - Billing/Accounting: cannot be scalable done in ASIC,
    - Fine-grained and specialized code: ASIC doesn't allow sufficient programmability for arbitrary code,
    - Interface and API for deployment: Automated attestation and provisioning.
- Measure / Benchmark VNFs
  - Dimensions to analyze and compare
    - Describe and quantify input of: SFC, operating system and hardware stack, Versioning, collateral workload, etc. etc.
    - Versioning: code release, configuration change.
    - Hardware specifications: x86 vs. FPGA vs GPU with performance guarantees (consistent / constant performance behavior).
    - Guarantees under normal and adversarial workloads: e.g. physically co-located VNFs competing for I/O or trashing memory caches.
    - Cost vs. performance trade-off: Cost performance characteristics or guarantees of a VNF.
  - Who uses the benchmarks? How much control does the benchmark have?
    - An entity using a third-party cloud/cloudlet?
    - An entity that controls code and infrastructure?
  - Open / General framework for testing (functional, regression, performance) during development. See Fd.io approach.
  - Linked-in versus Telco data centers.
  - Containers vs. VM: how do you reason about the guarantees.
  - Existing infrastructure support will impact how we design and decompose functions.
  - Definition of a VNF: Composed of several (one or more) virtualization units connected by an infrastructure network.
    - Think more broadly about not just a unit but an orchestration of VNFs.
    - SDN in the VNF and outside of the VNF.
    - VM in the data path – there are complex (legacy) constraints between the VM.
    - VM in the control space – multiple VMs for redundancy.
    - ETSI has the notion of VNFC which are components of the VNF – the potential unit.
- Ongoing / Related Work:
  - Gym [*Christian Rothenberg* & Ericsson Research][2]
  - vnfbench[3]
  - VBaaS[4]
  - Towards a Structured Approach to Developing Benchmarks for Virtual Network Functions.[5]
  - *Andreas Kassler*'s recent project (see Lightning talks).

---

2 http://materials.dagstuhl.de/files/17/17032/17032.ChristianEsteve%20Rothenberg1.Slides.pdf
3 https://tools.ietf.org/html/draft-rosa-bmwg-vnfbench-00
4 https://datatracker.ietf.org/doc/draft-rorosz-nfvrg-vbaas/
5 Jeremias Blendin, Yuriy Babenko, Dennis Kusidlo, Georg Schyguda and David Hausheer. In EWSDN 2016.

- Discussion:
  - *Dirk Kutscher:* The evolution of non-telco datacenters might be very relevant to the function placement / composition problem, e.g., the way that LinkedIn is running $change_me?
  - *Diego Lopez:* The VNF does not necessarily equal to a virtualization unit, e.g. a VNF can be composed out of more than one VM.
- To do:
  - Definition of a VNF highly unclear → a definition should be added to the report → discuss with *Diego Lopez*.
  - The VNF can contain forwarding, as well.
  - In ETSI, the term VNFC is used (c for component).
  - *Fabian Schneider:* It would be good to outline deployment options for VNF.
  - Container-based
  - In switches

5.5 NFV Use Cases

- Operator view
  - NFs in operator networks
  - Transparent middleboxes (e.g., proxy, application acceleration) are going to die – because all traffic is encrypted.
  - At the wired edge as a DSLAM and CGNAT.
  - In mobile networks there are eNodeBs, S-GWs (end of encrypted tunnels), and P-GWs.
  - Control boxes: accounting, radio service, monitoring.
  - We want to move all of these to software.
  - Vision of moving virtualized functions to the eNodeB and moving them to another eNodeB when the user moves.
  - NFV challenges: security, placement (includes hardware/software split), testing/visibility/verification, performance, economics?
  - We want to combine wired and mobile networks into one network.
- End users (and apps?) view
  - We want to insert VNF along the path.
  - This is being done a little by allowing insertion of, for example, virtual firewalls in the mobile edge cloud – only tenants with a business relationship with the provider are allowed to put VMs in the mobile edge cloud; mobile edge computing is like a CDN edge.
  - Not really cloud computing – call it “middleboxes-as-a-service”.
  - Service providers have same perspective.

6 Wrap-up

6.1 Research Areas and Questions

At the end of the seminar, we asked each seminar participant to identify research areas and open research questions on NFV and anonymously provide them on paper cards. We state the returned questions below.

- Impact of a microservice architecture on NFV performance, flexibility, resilience, etc.
- Metrics-Analytics-Policing-Control Loops: How much of network management can be automated and how can possible fatal clashes between control loop actions (e.g. one loop changes routing in one direction of another in the opposite direction) be identified and mitigated?
- Distributed Cloud, Edge Cloud: Does this have any impact on NFV deployment?
- How can HW capabilities (e.g. programmable hardware) be exposed by the hypervisor and used for NFV (placement)?
- How can a provider run untrusted NFs safely? How can a customer run NFs on an untrusted provider safely?
- How can an NFV platform provide real-time performance guarantees?
- How do you estimate the performance of a VNF in an early stage of the development?
- How to characterize/predict the performance of complex/dynamic systems with many VNFs.
- Placement optimization that considers the peculiar properties (e.g. the number of lookups, cache behavior). Optimized schedulers for this environment.
- How to include FPGA resources into NFV infrastructures while keeping elasticity and flexibility properties known from COTS processors?
- Networking vs. Commodity: What area should we match to?
- What sorts of applications would external parties (e.g. CDNs) like to deploy as network functions in carrier networks? One example might be a distributed denial of service mitigation function.
- How could carriers safely allow the deployment of such functions?
- What functionality could NFV provide that can't be easily provided any other way?
- Does it make sense to unify SDN with some sort of NFV control?
- Applicability of Microservice Architectures to Telco NFV.
- Better concepts for edge computing.
- Usable and Actionable SLAs for NFV.
- NFV introducing exponentially-growing heterogeneity may require to rethink current approaches to Data and Information Modelling.
- Methods to measure and guarantee consistent performance, including multi-dimensional KPIs.
- Computer-assisted de-composition & re-composition of network functions: Programming Language, Instantiation Optimization.
- How to decompose big, monolithic NFs into smaller microservices? Can we do this automatically? Creation of use cases beyond current ASIC/HW-centric NFs to drive NFV (open problem rather than specific question).
- How to enable inter-AS VNFs that span over administrative boundaries?
- Definition of metrics and performance evaluation principles for VNFs.
- How can you quantify the user-level impact of employing NFV? Performance impact? Privacy? Availability?
- How can you reason about the implication of configuration changes in a principled manner?
- How do we convert home network gateway devices, laptops, cell phones into infrastructures (Edge cloud) for supporting user / application-specific NFV?
- How do we perform proactive data plane verification in a network that uses traditional routing protocols?
- How do we build higher-fidelity – yet tractably analyzable – models of NFs for verification?

- How do we speed up the NF certification process performed by operators?
- How to unify a network containing thousands of stateful data plane nodes?
- Which new security threats arise with NFV?
- What can be offloaded to (programmable) hardware?
- Automatic leveraging of hardware and in-network capabilities for VNFs.
- YANG, TOSCA, ______ ? What comes next?[6]
- How to bring performance benchmarking results and placement/optimization together without getting killed by complexity; where is the right abstraction?
- How to design appropriate mechanisms to "check" decisions/output of machine learning models; how to do a quality assessment of the machine learning results on short time scales?
- What fundamental design differences in NFV architectures we would envisage if we assume the Datacenter network has 500 terabytes per second of bi-sectional bandwidth within a single datacenter (Google Prediction)?
- Making SDN and NFV convergence happen.
- A general theory (and practice!) of software-based networking.
- Data-driven network management.
- Security in all aspects.
- New business model
- Application of Machine Learning to function placement
- Techno-economic analysis
- Placement w.r.t. dynamics, functional decomposition
- Flexibility as a metric to analyze designs
- Verification

6.2 Lessons Learned, Surprises

We anonymously asked the seminar participants about their lessons learned during the seminar and state them below.

- The extent of the TUM work on EPC decomposition virtualization. They seem to have covered most of the important points.
- Significant disconnect between academia and industry.
- Control may or may not be different from management.
- How messy and complex the services are inside telco providers.
- There is no clear distinction in NFV: Middlebox vs. edge-cloud vs. 5G core virtualization.
- There is no clear consensus on the definition of a VNF.
- The distinction between conceptual and implementation aspects is fuzzy (e.g. VM/Container).
- That there is little excitement and novelty. Unfortunately.
- Pessimism about the impact that OpenFlow will have in the future (You can't buy OpenFlow switches)
- NFV isn't really about virtualization, it is about implementing network functionality in software – in a portable way.

[6] https://www.ietf.org/mail-archive/web/teas/current/msg01900.html

- SDN isn't about software, it is about separating the control plane from the dataplane and perhaps centralizing the control plane.
- Lack of common views and amount of disagreement.
- The amount of faith in Machine Learning / Artificial Intelligence to solve "untrackable" challenges in NFV/SDN.
- The definition of a VNF and NFV use cases is still a highly controversial topic that has not settled yet.
- NFV still is telco-centric: needs of service providers other than telcos do not seem to play a major role currently.
- The role of NFV in 5G networks.
- The plethora of open questions in NFV placement.
- There are research questions in NFV.
- Most researchers do not read architectural documents developed by SDOs in their domain. This leads to unnecessary discussions on terminology and common understanding of principles and concepts. *Point from the discussion: A literature research is not complete if you only consider academic papers.*
- Recent progresses in (control plane) verification.
- Foglets maybe make it to real deployments (just making fun...).
- American people do not say "allow to".
- Many people say there are no research challenges for NFV: No!!! When it comes to operational questions etc.
- "Intent" is not the powerful buzzword, I thought it was.
- Good old control plane protocols have still a lot to squeeze out when it comes to verification, control, management...
- People mistake OpenFlow to be equal to SDN!
- The control vs. management question is still not solved in industry.
- Confirmation of: SDN is not OpenFlow, which is dead.
- The NFV ecosystem (users, operators, vendors) needs further consideration.
- ETSI NFV standardization status (still fuzzy)

6.3 Controversial statements about NFV

In this session, we asked the 20 participants to write down controversial statements about NFV anonymously on blank cards. We collected them afterwards, read them aloud, and asked all participants to raise hands if they agree on the statement (👍) or if they disagree on it (👎).

"NFV is an attempt to re-invent Intelligent Networks."
👍 8 👎 8

"Will the network management overhead kill network virtualization and NFV?"
👍 0 👎 21

"SDN will never make it to the CORE of the Internet backbone"
👍 3 👎 18

"NFV will never make it to the CORE of the Internet backbone"
👍 13 👎 8

"There is no technology which is generic enough to solve different use-cases. We will end up in use-case specific solutions! Even P4 will (probably) be used to the interconnection DL/WAN. SDN is "only" used for TE in WAN (or what might be something like SDN)."
👍 1 👎 20

"SDN and NFV is both all about network management and not about control! (And network management is not well researched)"
👍 3 👎 12

"There is no research challenge in NFV"
👍 0 👎 21

"Data-driven (ML)-based network management will render networks infeasible to debug"
👍 7 👎 10

"Because SW is slower than optimized hardware, we need to throw much more hardware at the problem to meet current performance criteria."
👍 11 👎 6

"Like SDN, in practice, NFV apps will run only on end hosts, and not on network hardware/appliances deployed in networks"
👍 0 👎 21

"NFV will enable more smaller-size companies to enter NF markets (e.g. providing CDNs)"
👍 15 👎 5

"We (terribly) lack good software engineering principles in the application of SDN/NFV"
👍 15 👎 3

"Regulation will anyway restrict SDN & NFV to the enterprise networks"
👍 0 👎 21

"There has been no fundamental and deployed innovation in network management for the last"
👍 0 👎 0

"SDN/NFV as a research area is slowly but surely dying"
👍 5 👎 13

"NFV will increase costs due to Software maintenance and make misconfiguration even easier than today"
👍 8 👎 11

"Heterogeneity of deployment options will result in unmanageable systems"
👍 5 👎 15

"OpenFlow is no longer required!"
👍 13 👎 7

"NFV won't significantly lower costs for operators – they will need to structurally reform or will die"
👍 13 👎 5

"NFV research is mostly dead"
👍 0 👎 0

"NFV will never see widespread real-world deployment"
👍 0 👎 21

"There will only be at most 5 different useful network functions (i.e. there is no need for an NFV app store"
👍 13 👎 6

"Edge clouds will never provide a useful service normal people care about"
👍 1 👎 20

"Only NFs without cross-flow state can be dynamically scaled, because state sharing is prohibitively expensive"
👍 5 👎 12

"SDN will make it, but not with OpenFlow (how we know it)"
👍 18 👎 2

"NFV is going to converge with SDN"
👍 **9** 👎 **7**

"High-throughput, ASIC-based network appliances are not going to be replaced by NFV."
👍 **17** 👎 **2**

"NFV will have its first widespread deployment in the 5G RAN and mobile core"
👍 **12** 👎 **1**

"OpenFlow will remain a niche protocol and will never achieve its initial promise to replace existing protocols due to the lack of suitable hardware available in a timely fashion."
👍 **21** 👎 **0**

"Constructing real networks based on NFV will require a large amount of system integration (i.e. glue code), and until that is sorted through, networks based on NFV will be more expensive than traditional networks, but afterwards will be much cheaper."
👍 **14** 👎 **1**

"The impact of open source on standardization will be disruptive and there will be less SDOs in the future."
👍 **9** 👎 **3**

"NFV will lower the innovation, because vendors cannot sell boxes anymore, so there is no incentive to develop novel functions."
👍 **0** 👎 **21**

"NFV allows to speed up NF performance."
👍 **10** 👎 **6**

"There is only one way to meet the required NF performance: Put the NF in the switch!"
👍 **1** 👎 **20**

6.4 Wrap-up Notes

Author: Leonhard Nobach

- *Orchestration* has been identified as the current "hype buzzword".
- When NFV is solved, what will be the next hype? For example how to automate network management? How to avoid feedback loops (critical)? How to maximize automation? How to bring data-driven network management to the masses? Still, no buzzword exists for these questions.

- Other problems to solve are *intent management* (on-going work in the ONF[7]), and occurring conflicts.
- There might be a lack of solving "new" problems with NFV, i.e. a lack of innovation.

A very important outcome was that there is still no consensus on a very simple question: the definition of a network function (NF). Existing definitions have been discussed in a recent survey [1]: An NF can be considered as a logical network building block doing tasks which go beyond (SDN) forwarding. Taking such a definition, Layer 3 forwarding (involving more complex decisions, exchanging data link addresses and decrementing time-to-live counters) can be considered as a grey area, while any *stateful* processing (stateful firewalls, dynamic NAT) could be clearly considered as an NF. This does not exclude the possibility that parts of these NFs are implemented on an SDN data plane (for example the NF implementation dynamically adding or removing flows for established NAT or firewall sessions via OpenFlow).

According to this definition, a (V)NF can be implemented in one or multiple VMs (VNF instances), in the data plane itself, on reconfigurable, hardware–accelerated devices (e.g. FPGAs), or in a distributed system comprising a combination of these components.

References

1 Leonhard Nobach, Oliver Hohlfeld, David Hausheer: New Kid on the Block: Network Functions Virtualization: From Big Boxes to Carrier Clouds (Editorial). In: ACM Computer Communications Review (CCR), July 2016.

6.5 What comes after NFV?

Author: Oliver Hohlfeld

- Lots of pessimism in the research questions.
- Currently listed topics are more on the management side.
- Promise of SDN/NFV was to enable in the network. Where is this happening? *Comment James Kempf:* SDWAN.
- Don't apply verification to OpenFlow control plane aspects. Stop with OpenFlow. The important questions are in the management plane.
  - Verification for intent: ONF TIPI
  - Northbound interface of a controller.
  - Published four months ago, not yet accessible to research.
- *Remark Fabian Schneider:* Move the abstractions up the stack. We need abstractions going above OpenFlow and come with verification properties.
- *Comment Michael Scharf:* Scripting etc. needed way up in the management plane.
- To-do point: A literature research is not complete if you *only* consider academic papers.
- Microservice ecosystem (a.k.a. Active Networks)
- HW/SW Codesign (for flexibility)

[7] https://www.opennetworking.org/?p=1633&option=com_wordpress&Itemid=155

Participants

- Theo Benson
Duke University – Durham, US
- Christian Esteve Rothenberg
State University of Campinas, BR
- Aaron Gember-Jacobson
Colgate University – Hamilton, US
- Gabriela Gheorghe
PwC – Luxembourg, LU
- David Hausheer
TU Darmstadt, DE
- Oliver Hohlfeld
RWTH Aachen, DE
- Felipe Huici
NEC Laboratories Europe – Heidelberg, DE
- Andreas Kassler
Karlstad University, SE
- Wolfgang Kellerer
TU München, DE
- James Kempf
Ericsson – Santa Clara, US
- Dirk Kutscher
Huawei Technologies – München, DE
- Diego R. Lopez
Telefonica I+D – Seville, ES
- Bruce MacDowell Maggs
Duke University – Durham, US
- Oliver Michel
University of Colorado – Boulder, US
- Leonhard Nobach
TU Darmstadt, DE
- Costin Raiciu
University Politehnica of Bucharest, RO
- Fernando M. V. Ramos
University of Lisbon, PT
- Jan Rüth
RWTH Aachen, DE
- Michael Scharf
NOKIA – Stuttgart, DE
- Fabian Schneider
NEC Laboratories Europe – Heidelberg, DE
- Laurent Vanbever
ETH Zürich, CH
- Timothy Wood
George Washington University – Washington, DC, US
- Andreas Wundsam
Big Switch Networks – Santa Clara, US
- Thomas Zinner
Universität Würzburg, DE