NETWORK FUNCTIONS VIRTUALISATION – A CARRIER'S PERSPECTIVE

This article provides an insight into how Network Functions Virtualisation (NFV) has evolved over an approximately five-year timespan, what the current status looks like, and most importantly, what challenges remain to make this a mainstream technology. There is a particular focus on using innovative cross-industry collaborations to address many of these challenges.

A brief history

NFV stemmed from a number of industry collaborations involving focussed research into running "carrier-grade" telecoms network workloads on commoditised high-volume servers with Intel x86 multi-core processors. One such collaboration was formed by BT in 2011, and comprised Intel, HP, WindRiver, Verivue, F5 and Tail-F.

A very early and tangible outcome was a Proof-of-Concept (PoC) showing the achievable performance of a virtualised Broadband Remote Access Server as published in 2012 [1]. This PoC confirmed the huge potential of NFV as a means to unlock and de-couple the "high-value" software component of a network appliance from the underlying proprietary hardware, and thus leverage general-purpose servers to support "Virtualised Network Functions" (VNFs): crucially, the providers of the VNF software (routers, firewalls, load balancers, optimisation equipment, etc) could now be completely disaggregated from the

PAUL VEITCH, PETER WILLIS, PHILIP EARDLEY

State-of-play and future challenges

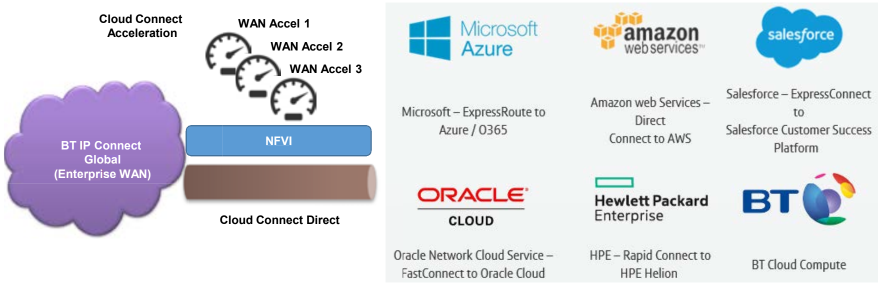

Figure 1: BT's Cloud Connect Acceleration Using NFV Infrastructure

hardware layer making for a much more flexible and vendor-portable network architecture [2].

In the autumn of 2012, a coalescence of network operators with a shared interest in NFV – instigated by BT and including Orange, Telefonica, Deutsche Telekom, Verizon and China Mobile – agreed to form a standards body under the auspices of the European Telecommunications Standards Institute (ETSI), and by early 2013 the inaugural ETSI Industry Specification Group (ISG) on NFV was launched[1]. The focus of the ETSI ISG has been to define a common and consistent set of architectural components, layers and interfaces for the NFV ecosystem stakeholders to adhere to: from the suppliers of the physical and virtualisation layers, through to suppliers of the VNFs and management components.

While some of the best brains in the network industry started defining what NFV should look like architecturally via standards fora including the ETSI ISG, operators were extremely active conducting practical lab experiments followed by actual pilots and trials, for specific NFV use cases. The early trials focussed not just on straightforward substitution of physical network functions with their virtual equivalent, but rather considered a much wider range of issues. These include aspects such as hypervisor[2] compatibility, performance characterisation, orchestration, and diagnostics.

The purpose of running such PoCs and trials has been to translate learning outcomes into mainstream design and development where "operating at scale" becomes a key factor. The early targets have been quick-win scenarios where there are obvious tangible business benefits to use VNFs running on standard servers rather than proprietary appliances, and, in the best-case scenarios have led to actual product development, as will be explained in the next section.

Current status

NFV underpinning real products

As of December 2015, BT Global Services had launched its first commercial product leveraging NFV. "Cloud Connect Acceleration"[3] provides optimisation and acceleration of data traffic sent between a large enterprise's global Wide Area Network (WAN) and third party cloud services, and is supported by virtualised Riverbed acceleration devices (called "Steelheads") running on standard servers. WAN accelerators reduce protocol chattiness and detect patterns of repetition in data packets, with smaller-sized tokens sent in place of repetitive (cached) data packets. This significantly reduces the number of packets sent over the WAN or viewed another way, accelerates the end-to-end data transfer. There were several compelling reasons to use the NFV approach for this particular use case:

- Given this was a new service capability, it made sense to host multiple VNFs on common NFV Infrastructure from the outset, rather than hosting multiple "per customer" physical appliances. This provides a much more agile, cost-effective and scalable means to manage per-customer instances of WAN acceleration than would be afforded with separate appliances.
- BT had performed significant testing of the virtualised Steelhead capability in its NFV lab, to ensure stable and predictable performance [3].
- The WAN acceleration capability is a component part of an end-to-end service to connect customer's existing WANs to third party cloud-based services (Figure 1), hence the strong alignment and synergy with cloud-based technology.

In addition to the already-launched Cloud Connect Acceleration, a suite of new services and capabilities are being launched by BT Global Services as part of the "Dynamic Network Services" programme, including enhancements to Cloud Connect products and a new "Agile CPE (Customer Premises Equipment)" capability which runs multiple VNFs on a single server. This is termed "universal" CPE and was one of the early use cases for NFV, mainly due to the obvious benefits of consolidating network functions like routers, firewalls, WAN optimisers onto a single box, but also due to dramatic improvements in service delivery and reduction in truck rolls associated with an enterprise network build-out.

NFV has been successfully leveraged by BT in certain other use cases, some of which are part of internal infrastructure, as well as underpinning services delivered to end customers:

- Use of a Metaswitch Session Border Controller in support of a "BT One Phone" product, enabling convergence of fixed, mobile and office phone systems.
- Use of Huawei Policy Control system as part of EE mobile core infrastructure.

[1] See: http://www.etsi.org/technologies-clusters/technologies/nfv

[2] The Hypervisor, sometimes called a Virtual Machine Monitor is the low-level programme used to host multiple guest operating systems, and manages access to server resources.

[3] See: https://www.globalservices.bt.com/uk/en/products/cloud-connectivity

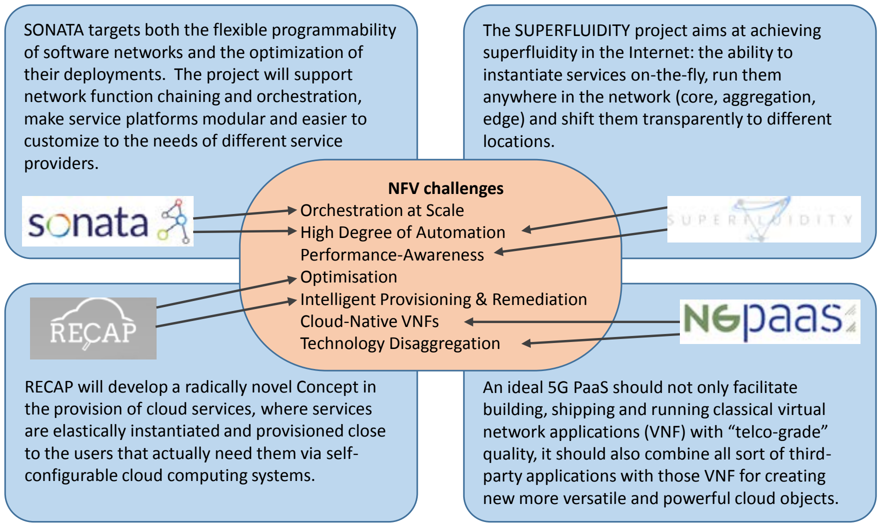

Figure 2: EU-funded projects addressing NFV challenges

NFV industry state-of-the-art

The previous section explained how NFV progressed from lab evaluations and trials, through to real-world deployments underpinning products and services. The industry itself has seen significant changes and developments in the space of five years. The ETSI ISG membership stands at over 300 companies including 38 service providers. A large proportion of those companies are established networking corporations or "Independent Software Vendors" who have a sizeable portfolio across a wide range of networking domains: virtual private networks, mobile (4G, emerging 5G), Content Delivery Networks (CDNs), broadband, Voice-over-IP (VoIP), test & diagnostics, quality monitoring, etc. The ecosystem also includes players who specialise in the Management and Orchestration (MANO) domain. This outcome truly reflects the unparalleled scale of the "shake-up" and momentum achieved by the NFV movement within the telecoms industry.

As well as ETSI, there are a growing number of open source bodies with a focus on NFV developments, some notable ones are as follows:

- Open Platform for NFV: whose goal is to "facilitate the development and evolution of NFV components across various open source ecosystems".
- Open Source MANO: aims to "develop an Open Source NFV Management and Orchestration (MANO) software stack aligned with ETSI NFV".
- Central Office Re-architected as a Data Centre: whose mission is to "combine NFV, SDN (Software Defined Networking), and the elasticity of commodity clouds to bring datacenter economics and cloud agility to the Telco Central Office".
- Intel Open Network Platform (ONP): defines itself as "a reference architecture that provides engineering guidance and ecosystem enablement support to encourage widespread adoption of SDN and NFV solutions in Telco, Enterprise and Cloud".
- Open Network Automation Platform: strives to develop a "comprehensive platform for real-time, policy-driven orchestration and automation of physical and virtual network functions that will enable software, network, IT and cloud providers and developers to rapidly create new services".

While formation of such bodies is for the large part a positive reflection of the collective desire for industrial cooperation and innovation in the NFV domain, there is some risk of competing initiatives trying to address the same problems in slightly different ways, thereby diluting the net benefits. From a carrier perspective, it can be a challenge to determine how much resource to allocate in active participation and/or tracking of such fora, especially when their mission statements sound similar.

One area where initiatives like OPNFV and Intel ONP have proved extremely beneficial is leveraging best-of-breed industry expertise in areas like performance optimisation. Servers can be set-up with "out-of-the-box" default settings, or they can be set-up with a wide array of potential tunings and optimisations to target more demanding "Telco-grade" workloads which require predictable performance in terms of latency, jitter and throughput. In conjunction with Intel and certain VNF vendors, BT has conducted a number of lab evaluations of optimisation technologies including Data Plane Development Kit (DPDK) and Cache Allocation Technology (CAT). These test scenarios validated the use of DPDK for low-latency applications (e.g. VoIP) [4], and how CAT can be used to mitigate against so-called "Noisy Neighbour" effects caused by resource-hungry VNFs consuming Last Level Cache resources on a central processor [5].

NFV has undoubtedly come a long way since 2012 and galvanised the telecoms industry into a new "software-defined" way of thinking. That said, has it achieved all it set out to? Table 1 compares the achievements and remaining gaps; each of the "Remaining Challenges" is intentionally set as a counter-point to the corresponding claimed "Success".

Collaborative innovation to address NFV challenges

Many of the gaps and challenges listed in Table 1 are being actively addressed in industry. BT for example is leveraging collaborations in a number of key EU-funded projects which will help drive solutions to some of the known problems. This is illustrated at a high-level in Figure 2, with additional detail on each project provided in the following sub-sections.

Sonata

The Sonata project[4] commenced in July 2015 and will complete in December 2017.

[4] See: http://sonata-nfv.eu/

Successes	Remaining challenges and gaps
Very wide range of server options (different cost points, form factors, etc) mostly based on Intel x86 but increasingly on alternatives like ARMv8.	The ARMv8 ecosystem support for mainstream VNF vendors, is still in its infancy.
Large number of vendors with VNFs.	Commercial VNF license models should be more attractive with a flexible range of options.
Large number of vendors with MANO solutions for NFV.	Certain vendor offerings are still regarded as turn-key solutions that only work with "vendor's own" VNFs, rather than being able to support a wide range of 3rd party VNFs on heterogeneous infrastructure: greater vendor/technology disaggregation is required.
Good portability of hypervisor support for VNFs.	Most or all VNFs are built as monolithic "virtual replicas" of their complete operating systems that run on hardware: new "Cloud Native" technologies for VNFs like containers[5] are relatively immature and not yet adopted widely by VNF vendors.
Good understanding of performance characterisation and available "fine-tuning" options of standard compute infrastructure to achieve predictable performance.	Applying performance-tuning settings in real-world deployments in an automated way, and including embedded awareness of the infrastructure set-up at the MANO layers.
Industry standards bodies (e.g. ETSI ISG) providing common focus for players in ecosystem.	Uncertainty remains as to what a truly "standards-compliant" NFV implementation looks like.
Initial deployments supporting real-world products (albeit relatively modest in scale in terms of managed instances).	The full potential of NFV where large numbers of VNFs can be auto-provisioned, configured, managed, logically-connected in "Service Chains", scaled up/down as needed and auto-remediated when failures occur, has yet to be unleashed.
Rich ecosystem in open source communities ensuring low-cost implementation options available.	Determining how and where to target effort and priority in a somewhat crowded space of similar-looking initiatives.
Good understanding of security challenges, with industry focus in ETSI ISG "SEC" Working Group.	Achieving consistent approach to addressing security issues in real-world, large-scale deployments, while also addressing emerging threats in future deployments such as 5G.

Table 1: NFV successes and remaining challenges

This consortium comprises a mix of industry partners (including NEC, Telefonica and BT) and academic institutions (including University College London, and the University of Paderborn).

The top-level goal of Sonata is to develop open-source tools and capabilities for orchestrating a range of VNF types across a wide spectrum of service provider use cases, while accelerating the development and testing of NFV-based services. It has achieved this through the specification and development of the Sonata "Service Platform" (Figure 3). Some key characteristics include:

- An extendable and modular "plug-in" architecture implemented using micro-services which allows the owner to make alterations such as addition and removal of plug-ins on the fly.
- The ability to abstract the underlying infrastructure meaning more than one type of Virtual Infrastructure Manager can be orchestrated, e.g. Openstack, Vmware, etc.
- A Software Development Kit toolkit: to ease network service development by service developers.

One of BT's contributions to this project has been to specify the orchestration requirements for the virtual CDN use case, and how two or more service providers can cooperate when using NFV.

Superfluidity

The Superfluidity project[6] commenced in July 2015 and will complete in March 2018.

[5] Containers are isolated systems/applications sharing the Linux kernel. They leverage operating system-level virtualisation rather than hardware-based virtualisation enabling more resource-efficient VNFs.

[6] See: https://5g-ppp.eu/superfluidity/

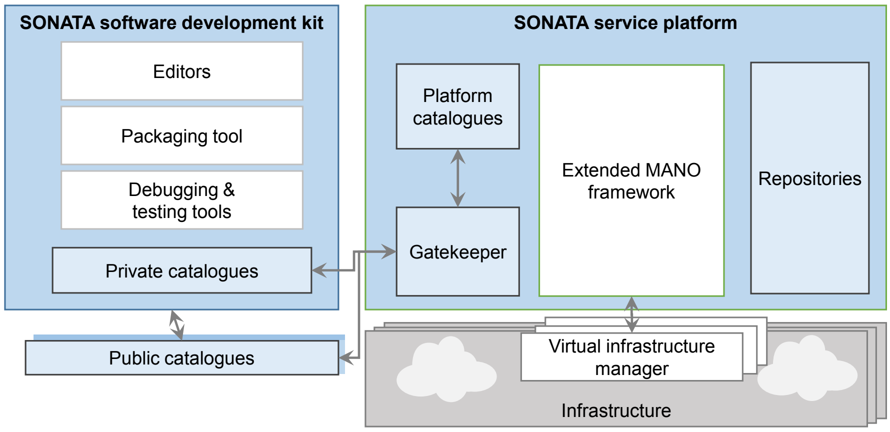

Figure 3: Overview of SONATA service platform

This consortium comprises a mix of industry partners (including Citrix, BT and Redhat) and academic institutions (including University of Liege and the Technical University of Dresden).

The main themes of this project have been large-scale management and ways to exploit heterogeneity of orchestration systems and underlying NFV infrastructure implementations. In terms of heterogeneity, this is increasingly important due to the wide range of potential NFV implementations that will be required to support future services. For example, a major consideration is the impact that 5G converged networks will have. Fixed and mobile technologies will be increasingly unified and a very diverse range of services linked to emerging vertical industries like automotive, Internet-of-Things, smart energy, agriculture and healthcare will be supported on a common infrastructure [6].

It is generally accepted that technologies including NFV will be essential tools to meet the diverse and elastic demands of future 5G capabilities [7]. However, rather than a "one-size-fits-all" approach, different "flavours" of NFV implementation will co-exist. Alternative virtualisation technologies for VNFs, including classic Virtual Machines, could sit alongside smaller footprint VNFs (e.g. based on Containers), with the underlying physical infrastructure being made up of different hardware types, from basic un-optimised compute, through to performance-optimised compute (e.g. using DPDK, Cache Allocation technology, etc [4,5]), and even specialised hardware such as Field Programmable Gate Arrays (FPGAs). The Superfluidity project has made significant progress in defining a fully flexible architecture that meets many of these requirements; throughout the course of this project, BT has hosted a testbed at the Adastral Park campus for all the partners.

RECAP

The RECAP project[7] commenced in December 2016 and will run until December 2019. The full name of the project is "Reliable Capacity Provisioning and enhanced Remediation for Distributed Cloud Applications". As well as BT, this consortium comprises a mix of industry partners (including Intel and Tieto) and academic institutions (including Ulm, Dublin City University and Umea Universities).

This project is applying the concept of machine learning to the provisioning and remediation of NFV and Cloud infrastructures, thus applying a high level of computational intelligence to process large quantities of data gathered from network

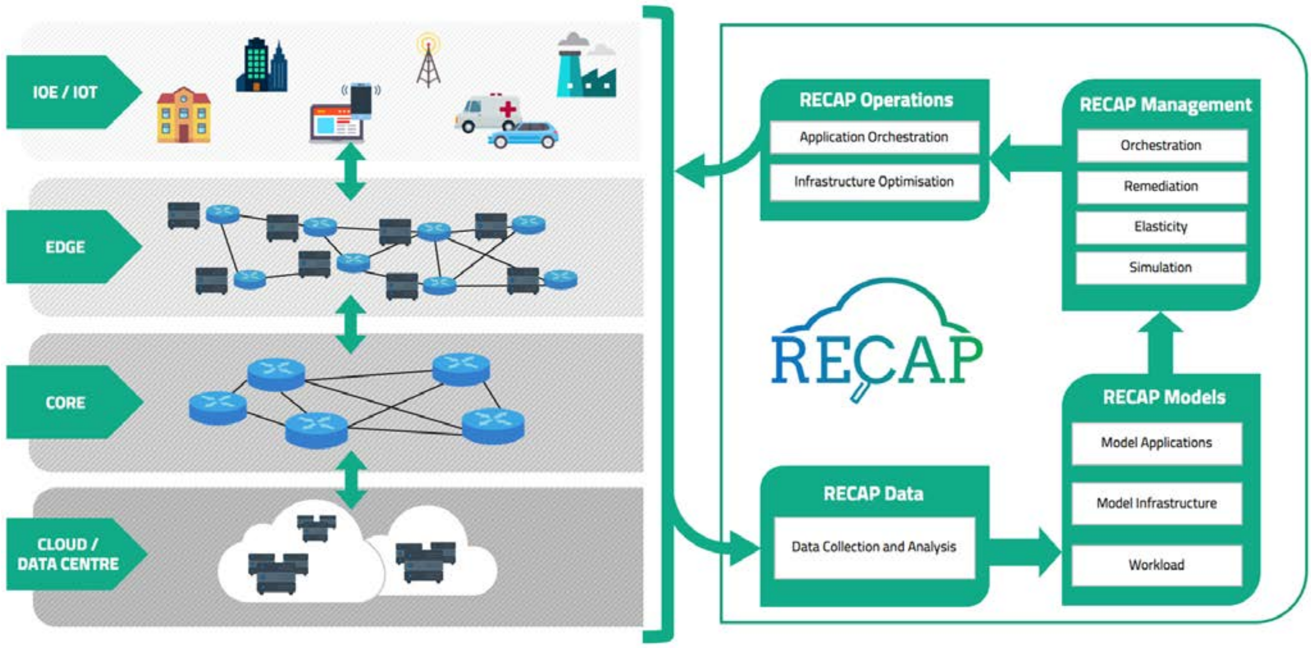

Figure 4: RECAP top-level objectives

[7] https://recap-project.eu/about/

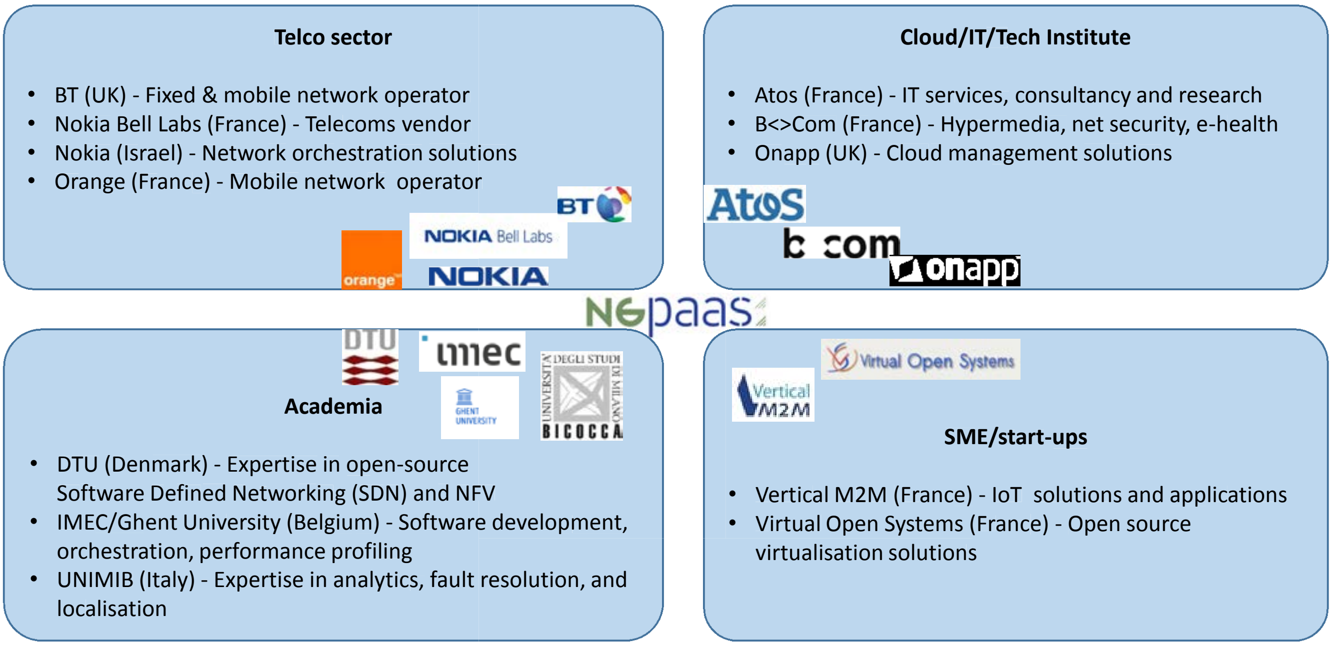

Figure 5: Diverse ecosystem of NGPaaS partners

and application telemetry. The use of machine learning techniques should enable much greater levels of predictability of network and application behaviour. It will achieve this by simulating the effects of different decisions related to application and component placement (e.g. specific VNFs could potentially be hosted at large number of possible distributed compute locations), with a view to optimising the performance and capacity utilisation. Figure 4 shows the RECAP objectives in a graphical format.

BT's contributions to RECAP include expertise on distributed NFV, network management, virtual CDN and Cloud Connect Application use cases as well as hosting part of the RECAP test lab.

NGPaaS

The Next Generation Platform-as-a-Service (NGPaaS) project[8] commenced in June 2017 and will complete in June 2019. Figure 5 uses a simple graphic to illustrate the partners involved in this particular project, which is also a useful way to demonstrate how collaborative projects of this nature comprise best-of-breed expertise from a range of industry sectors and academia.

NGPaaS has some ambitious goals to enable customisation of platforms used to support a range of telco-grade VNFs and other third party cloud objects in a unified manner. In simple terms this will require the development of VNFs which are more "cloud-native" in their build and implementation, but at the same time ensuring "telco-grade" aspects are facilitated in a completely disaggregated fashion. This should result in avoidance of technology lock-in at all levels of the stack, so for example at the physical level, ARMv8, Intel/AMD x86 and FPGA options are equally viable. Furthermore, the level of customisation available to "build-to-order" should extend to all the potential vertical market segments that will exist in future 5G converged infrastructures: mobile, fixed and a range of Internet of Things scenarios including smart cities, energy, agriculture, etc.

NGPaaS will showcase these new developments by adopting a continuous development and integration (i.e. "Devops") lifecycle, harnessed not just within, but across industry verticals. This new model is referred to as "Dev-for-Operations" and will be a key challenge of NGPaaS.

BT's contributions to NGPaaS include expertise on NFV use cases, architecture, and network management, as well as hosting part of the NGPaaS test lab at the Adastral Park campus for all the partners.

AUTHORS' CONCLUSIONS

NFV gained significant traction and momentum thanks to successful cross-industry collaborations. As well as leading to the first real-world trials and product-driven deployments of NFV, this resulted in an unprecedented growth in one of the largest ETSI ISGs of its kind in terms of contributors and participants, as well as myriad of open source fora and bodies.

In the same spirit, many of the challenges that currently stand between NFV being a technology used in a fairly ad hoc and piecemeal fashion, to something that is all-pervasive in next-generation networks and platforms such as 5G, can only be fully addressed by maximising collaborative innovation between best-of-breed industrial and academic partnerships. This article highlighted a number of key EU-funded collaborative projects that strive to address many of these remaining challenges.

[8] http://ngpaas.eu/

REFERENCES

1. Willis, P. Virtualising Network Services in the Cloud. Carrier Cloud Summit, Paris, Jun 2012
2. Network Functions Virtualisation. Introductory White Paper, Octr 2012. Available at: https://portal.etsi.org/nfv/nfv_white_paper.pdf
3. Veitch, P., Mcgrath, M., and Bayon,V. An Instrumentation and Analytics Framework for Optimal and Robust NFV Deployment. IEEE Communications Magazine, Feb 2015
4. Veitch, P., Long, T., and Hutchison, P. NFV Performance Optimization for Virtualized Customer Premises Equipment. Intel Developer Zone, Aug 2016. Available at: https://software.intel.com/en-us/articles/nfv-performance-optimization-for-vcpe
5. Veitch, P., Curley, E., and Kantecki, T. Performance Evaluation of Cache Allocation Technology for NFV Noisy Neighbor Mitigation. IEEE Netsoft 2017, Bologna, Jul 2017
6. View on 5G Architecture. White Paper by the 5G PPP Architecture Working Group, Jul 2016. Available at: https://5g-ppp.eu/wp-content/uploads/2014/02/5G-PPP-5G-Architecture-WP-July-2016.pdf
7. Network Functions Virtualisation: Network Operator Perspectives on NFV priorities for 5G. ETSI White paper, Feb 2017. Available at: https://portal.etsi.org/NFV/NFV_White_Paper_5G.pdf

ABOUT THE AUTHORS

Paul Veitch holds M.Eng. and Ph.D. degrees from the University of Strathclyde, Glasgow. He joined BT in 1996, and worked on core transmission, multi-service platforms, and 3G mobile infrastructure design, before moving to Verizon Business (UUNET) in 2000. Paul returned to BT in 2003, and was infrastructure design authority for IP VPN and BT consumer networks. In his current role he leads NFV trials with a focus on business down-streaming, and manages BT's contribution to NGPaaS.

Peter Willis researches the future of networks. He has been researching and developing Network Functions Virtualisation since 2011, he published the first carrier NFV testing results in June 2012 and is co-inventor of the term "NFV". Peter is currently leading BT's research to improve NFV technology and its management. Peter previously worked on the development of PBB-TE, BT's 21st Century Network Architecture and BT's Internet service.

Philip Eardley works in the Converged Networks team within BT Research & Innovation. He researches the future converged digital infrastructure, which will incorporate technologies such as 5G, virtualisation and autonomics. He led the EC collaborative projects Trilogy ('Architecting the future Internet') and Leone ('From global measurements to local management'). He chairs the IETF's Multipath TCP working group. He co-authored the book IP for 3G.

ITP AUTHORS
Want to know more? To contact the authors email your name, company name and email address to thejournal@theitp.org

ABBREVIATIONS

CAT	Cache Allocation Technology
CDN	Content Delivery Networks
CPE	Customer Premises Equipment
DPDK	Data Plane Development Kit
ETSI	European Telecommunications Standards Institute
FPGA	Field Programmable Gate Array
ISG	Industry Specification Group
MANO	Management and Orchestration
NGPaaS	Next Generation Platform-as-a-Service
NFV	Network Functions Virtualisation
ONP	Open Network Platform
PoC	Proof-of-Concept
SDN	Software Defined Networking
VNF	Virtualised Network Function
VoIP	Voice-over-IP
WAN	Wide Area Network

ACKNOWLEDGEMENTS

The research described above in Sonata has been partly funded by the European Commission H2020 5G-PPP project SONATA (671517).

The work described above in Superfluidity received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 671566.

The work described above in RECAP received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 732667 (RECAP).

The work described above in NGPaaS received funding from the European Commission under the Horizon 2020 and 5G-PPP Phase2 programmes (Grant Agreement No. 761 557).