Ericsson Review

The communications technology journal since 1924

2013 • 2

Software-defined networking: the service provider perspective

February 21, 2013

ERICSSON

Software-defined networking: the service provider perspective

An architecture based on SDN techniques gives operators greater freedom to balance operational and business parameters, such as network resilience, service performance and QoE against opex and capex.

ATTILA TAKACS, ELISA BELLAGAMBA, AND JOE WILKE

The traditional way of describing network architecture and how a network behaves is through the fixed designs and behaviors of its various elements. The concept of software-defined networking (SDN) describes networks and how they behave in a more flexible way – through software tools that describe network elements in terms of programmable network states.

The concept is based on split architecture, which separates forwarding functions from control functions. This decoupling removes some of the complexity from network management, providing operators with greater flexibility to make changes.

Ericsson's approach to SDN goes beyond the data center addressing issues in the service-provider environment. In short Ericsson's approach is Service Provider SDN. The concept aims to extend virtualization and OpenFlow – an emerging protocol for communication between the control and data planes in an SDN architecture – with three additional key enablers: integrated network control; orchestrated network and cloud management; and service exposure.

There is no denying that networks are becoming increasingly complex. More and more functionality is being integrated into each network element, and more and more network elements are needed to support evolving service requirements – especially to support rising capacity needs, which are doubling every year[1]. One of the root causes of network complexity lies in the traditional way technology has developed. The design of network elements, such as routers and switches, has traditionally been closed; they tend to have their own management systems with vertically integrated forwarding and control components, and often use proprietary interfaces and features. The goal of network management is, however, to ensure that the entire network behaves as desired – an objective that is much more important than the capabilities of any single network element. In fact, implementing end-to-end networking is an important mission for most operators, and having to configure individual network elements simply creates an unwanted overhead.

Network-wide programmability – the capability to change the behavior of the network as a whole – greatly simplifies the management of networks. And the purpose of SDN is exactly that: to be able to modify the behavior of entire networks in a controlled manner.

The tradition of slow innovation in networking needs to be broken if networks are to meet the increased demand for transport and processing capacity. By integrating recent technological advances and introducing network-wide abstractions, SDN does just that. It is an evolutionary step in networking.

Telephony has undergone similar architectural transitions in the past. One such evolution took place when a clear separation between the functions

BOX A Terms and abbreviations

API	application programming interface	ONF	Open Networking Foundation
ARPU	average revenue per user	OSS/BSS	operations and business support systems
CLI	command-line interface	PE	provider edge device
DPI	deep packet inspection	PGW	Packet Data Network Gateway
GMPLS	generalized multi-protocol label switching	RG	residential gateway
L2	Layer 2	RWA	routing and wavelength assignment
L3	Layer 3	SDN	software-defined networking
L2-L4	Layers 2-4	SGW	Service Gateway
M-MGW	Mobile Media Gateway	SLA	Service Level Agreement
MME	Mobility Management Entity	VHG	virtual home gateway
MSC-S	Mobile Switching Center Server	VoIP	voice over IP
NAT	Network Address Translation	WAN	wide area network
NMS	network management system		

FIGURE 1 **Service Provider SDN – components and promise**

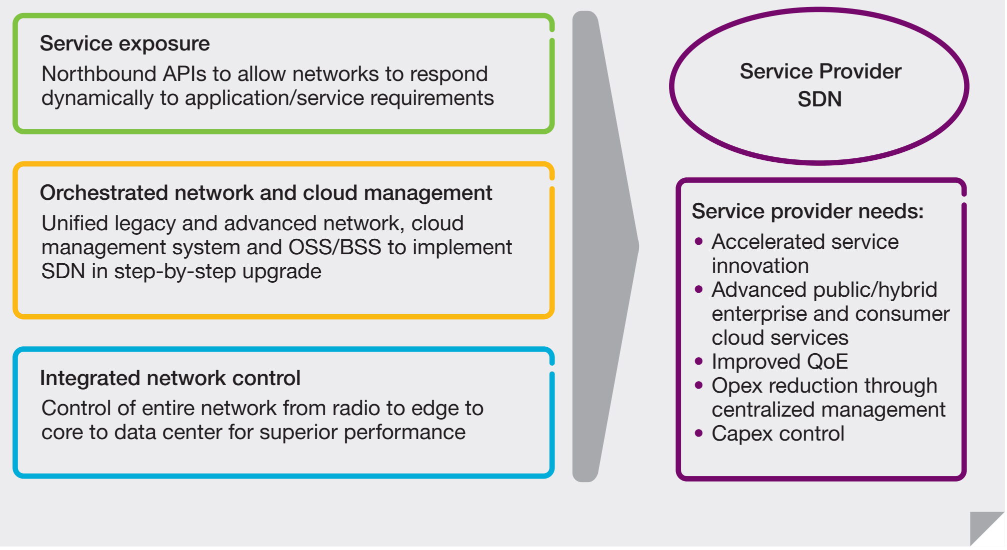

of the data plane (including SGW, PGW and M-MGW) and the control plane (including MME and MSC-S) was introduced. Now SDN has brought the concept of split architecture to networking.

As the business case proves, over the next two to five years, SDN technology will be deployed in networks worldwide. At the same time, the need to maintain traditional operational principles and ensure interoperability between SDN and more traditional networking components will remain. In the future, SDN will help operators to manage scale, reduce costs and create additional revenue streams.

Standardization

The goal of the Open Networking Foundation (ONF), which was established in 2011, is to expedite the standardization of the key SDN interfaces. Today, the work being conducted by ONF focuses on the continued evolution of the OpenFlow protocol. ONF has recently established the Architecture and Framework Working Group, whose goal is to specify the overall architecture of SDN. The work carried out by this group will guide future standardization efforts based on strategic use cases, requirements for data centers and carrier networks, the main interfaces, and their roles in the architecture. Ericsson is actively driving the work of this group, cooperating with other organizations to promote the evolution of OpenFlow and support an open-source implementation of the most recent specifications.

Other standardization organizations, most notably the IETF, have recently begun to extend their specifications to support SDN principles. In IETF, the Interface to the Routing System (i2rs) WG and the recent activity of the Path Computation Element (PCE) WG will result in standardized ways to improve flexibility in changing how IP/MPLS networks behave. This is achieved through the introduction of new interfaces to distributed protocols running in the network, and mechanisms to adapt network behavior dynamically to application requirements.

In addition to standardization organizations a multitude of active communities and open-source initiatives, such as OpenStack, are getting involved in the specification of various SDN tools, working on maturing the networking aspect of virtualization.

Architectural vision

Split architecture – the decoupling of control functions from the physical devices they govern – is fundamental to the concept of SDN. In split-architecture networks, the process of forwarding in the data plane is separated from the controller that governs forwarding in the control plane. In this way, data-plane and control-plane functions can be designed, dimensioned and optimized separately, allowing capabilities from the underlying hardware to be exposed independently.

This ability to separate control and forwarding simplifies the development and deployment of new mechanisms, and network behavior becomes easier to manage, reprogram and extend. Deploying a split architecture, however, does not remove the need for high-availability software and hardware components, as networks continue to meet stringent carrier requirements. However, the decoupling approach to architecture will help rationalize the network, making it easier to introduce new functions and capabilities. The ultimate goal of the SDN architecture is to allow services and applications to issue requests to the network dynamically, avoiding or reducing the need for human intervention to create new services. This, compared to today's practices, will reduce the time to market of new services and applications.

The OpenFlow protocol[2] is supporting the separation of data and control planes and allowing the path of packets through the network to be software determined. This protocol provides a simple abstraction view of networking equipment to the control layer. Split architecture makes virtualization of networking resources easier, and the control plane can provide virtual views of the network for different applications and tenants.

Ericsson has worked together with service providers to understand their needs for SDN both in terms of reducing costs and creating new revenue opportunities. Based on these discussions, Ericsson has expanded the industry definition of SDN and customized it to fit the needs of operators.

Ericsson's hybrid approach – Service Provider SDN – extends industry definitions including virtualization and OpenFlow with three additional key enablers:

- integrated network control;
- orchestrated network and cloud management; and
- service exposure.

FIGURE 2 Service Provider SDN – architectural vision

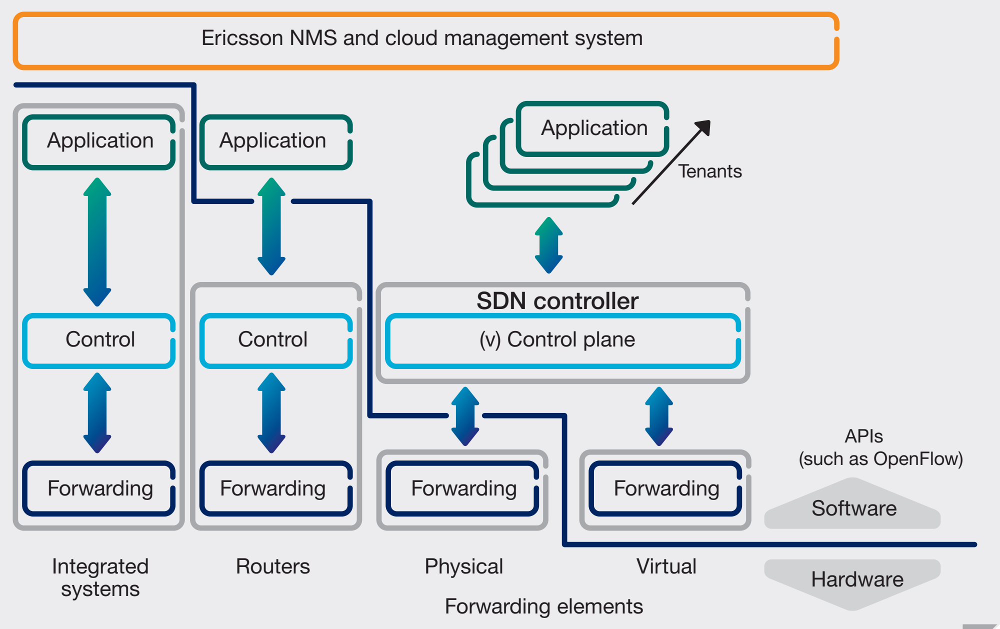

Integrated network control

Service providers will use SDN across the network from access, to edge to core and all the way into the data center. With integrated network control, operators can use their network features, including QoS, edge functions and real-time activity indicators to deliver superior user experience.

Orchestrated network and cloud management

Service Provider SDN will integrate and unify legacy network management systems with new control systems as well as with OSS/BSS. The platform for integrated orchestration supports end-to-end network solutions ranging from access over aggregation to edge functions as well as the data centers used to deliver telco and enterprise applications and services.

Service exposure

Northbound APIs expose the orchestration platform to key network and subscriber applications and services. Together, the APIs and platforms allow application developers to maximize network capabilities without requiring intimate knowledge of their topology or functions.

By expanding the perspective of SDN to include these three elements, service providers can evolve their existing network to the new architecture to improve the experience of their customers. Implementing Service Provider SDN should remove the dumb pipe label giving operators an advantage over competitors that do not own networks.

Network virtualization

One of the benefits of Service Provider SDN, especially from a network-spanning perspective is network virtualization. Through virtualization, logical abstractions of a network can be exposed instead of a direct representation of the physical network. Virtualization allows logical topologies to be created, as well as providing a way to abstract hardware and software components from the underlying network elements, thereby separating control from forwarding capabilities and supporting the centralization of control.

Unified orchestration platforms support network programming at the highest layer as programming instructions flow through the control hierarchy – potentially all the way down to granular changes in flow paths at the forwarding plane level. Adding northbound APIs into this unified orchestration layer provides the necessary support for applications and tenants to trigger automatic changes in the network, ensuring optimal QoS and guaranteeing SLAs.

Unified and centralized orchestration platforms greatly simplify the process of configuring, provisioning and managing complex service networks. Instead of having to tweak hundreds of distributed control nodes using fairly complex CLI programming, operations staff can use simple intuitive programming interfaces to quickly adjust network configurations and create new services. By accelerating the process of service innovation, Service Provider SDN will lead to increased market share and service ARPU, creating significant revenue growth and possibly reducing annual churn rates.

A high-level network architecture that supports the Service Provider SDN vision is illustrated in Figure 2. Service-provider networks will combine distributed control plane nodes (traditional routers and appliances), and data-plane elements that are governed by centralized elements – SDN controllers – in the control plane. Consequently, to make service-provider networks programmable, distributed and centralized control-plane components must be exposed to a unified orchestration platform. In addition, the key elements of the orchestration platform and the control plane need to be exposed to network and subscriber applications/services.

Application examples

The major use cases or applications of SDN in service-provider networks are summarized in Figure 3. Of these applications, the data center was the first to make use of SDN. Ericsson's approach to this is described in several articles published in Ericsson Business Review[3] and in Ericsson Review[4,5].

SDN can be applied in the aggregation network to support sophisticated virtualization and to simplify the configuration and operation of this network segment. Ericsson has developed a proof-of-concept system in cooperation with Tier 1 operators to evaluate the applicability of SDN to aggregation networks. This work has been carried out as part of the European Commission's Seventh Framework Programme[6].

Ericsson and Telstra have jointly developed a service-chaining prototype that leverages SDN technologies to enhance granularity and dynamicity of service creation. It also highlights how SDN can simplify network provisioning and improve resource utilization efficiency.

Packet-optical integration is a popular topic of debate, out of which several different approaches are emerging. Split architecture provides a simple way to coordinate packet and optical networking; and so SDN, enhanced with features such as routing and wavelength assignment (RWA) and optical impairments management will be a natural fit for packet-optical integration.

Ericsson has started to develop solutions to virtualize the home gateway. Virtualization reduces the complexity of the home gateway by moving most of the sophisticated functions into the network and, as a result, operators can prolong the home gateway refreshment cycle, cut maintenance costs and accelerate time to market for new services.

Virtualization of aggregation networks

The characteristics shared by aggregation and mobile-backhaul networks are a large number of nodes and relatively static tunnels – that provide traffic grooming for many flows. These networks are also known for their stringent requirements with regard to reliability and short recovery times. Besides L2 technologies IP and IP/MPLS is making an entrance as a generic backhaul solution. From an operational point of view, despite the availability of the distributed control plane technology, this network segment is usually configured statically through a centralized management system, with a point of touch to every network element. This makes the introduction of a centralized SDN controller straightforward for backhaul solutions.

A control element hosted on a telecom-grade server platform or on an edge router provides the operator with an interface that has the same look and feel as a single traditional router. The difference between operating an aggregation SDN network and a traditional network lies in the number of touch points required to provision and operate the domain. In the case of SDN, only a few points are needed to control the connectivity for the entire network. Consider, for example, an access/aggregation domain with hundreds or even thousands of nodes running distributed IGP routing protocols and the Label Distribution Protocol (LDP) to configure MPLS forwarding. In this case, SDN principles can be applied to simplify and increase the scalability of provisioning and operating of such a network by pulling together the configuration of the whole network into just a few control points.

The control element treats the underlying forwarding elements as remote line cards of the same system and, more specifically, controls their flow entries through the OpenFlow protocol. With this approach, any kind of connectivity model is feasible regardless of whether the forwarding node is L2 or L3 as, from a pure forwarding point of view, the same model is used in both flows. At the same time, network resilience at the transport level can be implemented by adding protection mechanisms to the data path. The SDN controller can pre-compute and pre-install backup routes and then protection switching is handled by the network elements for fast failover. Alternatively, the SDN controller can reroute around failures, in case multiple failures occur, or in scenarios that have less stringent recovery requirements.

From the outside, the entire network segment appears to be one big PE router and, for this reason, neighboring network elements of the SDN-controlled area cannot tell the difference (from a protocol point of view) between it and a traditional network. The network controller handles the interfacing process with legacy systems for connection setup. Additional information on this point can be found in a presentation on the Virtual Network System[7].

Dynamic service-chaining

For inline services, such as DPI, firewalls (FWs), and Network Address Translation (NAT), operators use different middleboxes or appliances to manage subscriber traffic. Inline services can be hosted on dedicated physical hardware, or on virtual machines. Service chaining is required to route certain subscriber traffic through more than one such service. There are still no protocols or tools available for operators to perform flexible, dynamic traffic steering. Solutions currently available are either static or their flexibility is significantly limited by scalability inefficiencies.

FIGURE 3 Application examples

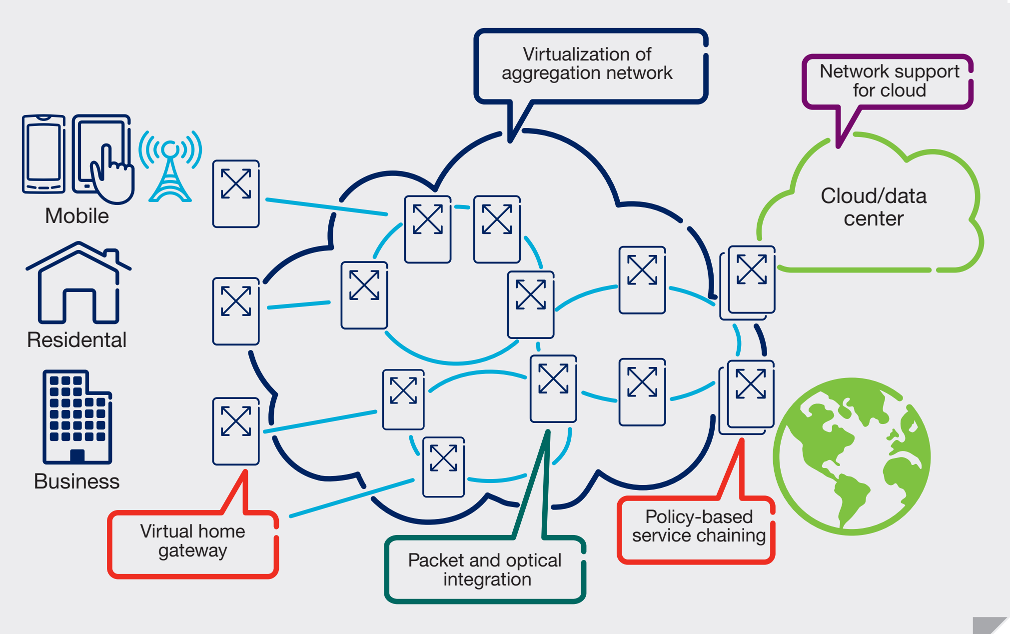

FIGURE 4 Service-chaining principles

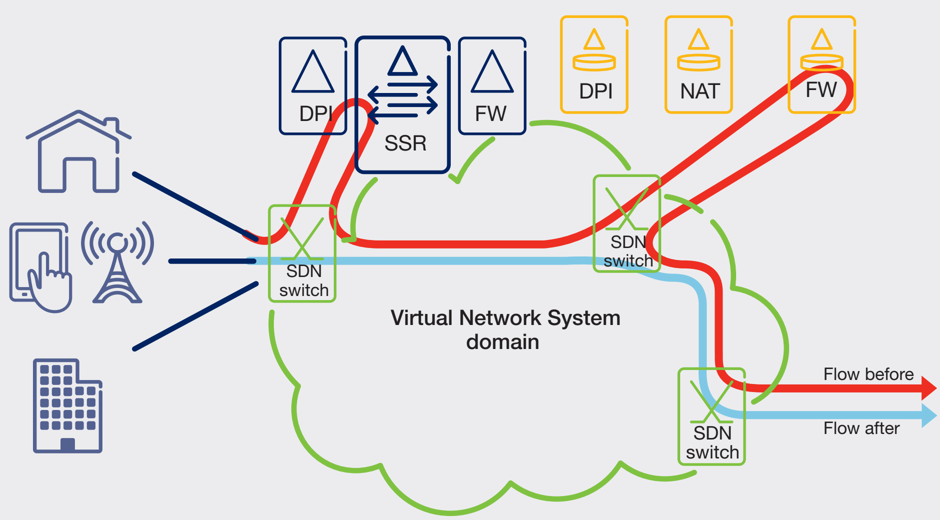

Given the rate of traffic growth, continued investment in capacity for inline services needs to be managed carefully. Dynamic service-chaining can optimize the use of extensive high-touch services by selectively steering traffic through specific services or bypassing them completely which, in turn, can result in capex savings owing to the avoidance of over-dimensioning.

Greater control over traffic and the use of subscriber-based selection of inline services can lead to the creation of new offerings and new ways to monetize networks. Dynamic service steering enables operators to offer subscribers access to products such as virus scanning, firewalls and content filters through an automatic selection and subscribe portal.

This concept of dynamic service chaining is built on SDN principles. Ericsson's proof-of-concept system uses a logically centralized OpenFlow-based controller to manage both switches and middleboxes. As well as the traditional 5-tuple, service chains can be differentiated on subscriber behavior, application, and the required service. Service paths are unidirectional; that is, different service paths can be specified for upstream and downstream traffic.

Traffic steering has two phases. The first classifies incoming packets and assigns a service path to them based on predefined policies. Packets are then forwarded to the next service, based on the current position in their assigned service path. No repeating classification is required; hence the solution is scalable.

The SDN controller sets up and reconfigures service chains flexibly to an extent that is not possible with today's solutions. The dynamic reconfiguration of service chains needs a mechanism to handle notifications sent from middleboxes to the controller. For example, the DPI engine notifies the controller that it has recognized a video flow. These notifications may be communicated using the extensibility features of the OpenFlow 1.x protocol.

Figure 4 summarizes service-chaining principles. The Virtual Network System (VNS) is a domain of the network where the control plane is centralized which excludes some of the traditional control agents. A simple API, such as OpenFlow, can be used to control the forwarding functionality of the network, and the VNS can create northbound interfaces and APIs to support creation of new features, such as service chaining, which allows traffic flows to be steered dynamically through services or parts thereof by programming forwarding elements.

The services provided by the network may reside on devices located in different parts of the network, as well as within an edge router – for example, on the service cards of Ericsson's Smart Services Router (SSR). Service chains are programmed into the network based on a combination of information elements from the different layers (L2-L4 and possibly higher). Based on operator policies, various services can be applied to traffic flows in the network.

For example, traffic may pass through DPI and FW functions, as illustrated by the red flow in Figure 4. However, once the type of the flow has been determined by the DPI function, the operator may decide to modify the services applied to it. For example, if the flow is an internet video stream, it may no longer need to pass the FW service, reducing load on it. Furthermore, after the service type has been detected, the subsequent packets of the same flow may no longer need to pass the DPI service either; hence the path of the flow can be updated – as indicated by the blue flow in Figure 4.

Packet-optical integration

The increased programmability that SDN enables creates an opportunity to address the challenges presented by packet optical networking. SDN can simplify multi-layer coordination and optimize resource allocation at each layer by redirecting traffic (such as VoIP, video and web) based on the specific requirements of the traffic and the best serving layer.

Instead of a layered set of separated media coordinated in a static manner, SDN could transform the packet-optical infrastructure to be more fluid, with a unified recovery approach and an allocation scheme based on real-time link utilization and traffic composition. The ONF still has some work to do to adapt OpenFlow to cope with optical constraints.

To speed up packet-optical integration, a hybrid architecture can be deployed where OpenFlow drives the packet domain, and the optical domain remains under the control of GMPLS. This approach utilizes the extensive optical capabilities of GMPLS and, therefore, instead working to extend OpenFlow with optical capabilities, it allows us to focus on the actual integration of optical and packet domains and applications that utilize the flexibility of a unified SDN controller.

Home gateway control

The concept of the virtual home gateway (VHG) introduces a new home-network architecture primarily driven by consideration to improve service delivery and management. The target architecture emerges by applying SDN capabilities between the residential gateway (RG) and the edge network – moving most of the gateway's functionalities into an embedded execution environment. Virtualizing the RG significantly reduces its complexity and provides the operator with greater granularity in remote-control management, which can be extended to every home device and appliance. As a result, operators can reduce their investments significantly by prolonging the RG refreshment cycle, cutting maintenance costs and accelerating time to market for new services.

The VHG concept allows operators to offer seamless and secure remote profile instantiation stretching the boundaries of a home network without compromising security. The concept provides the tools to configure and reconfigure middleboxes dynamically, so that communication between devices attached to different home networks can be established, and/or provide specific connectivity requirements for a third-party service provider – between, for example, a utility company and a particular device. By embedding SDN capabilities, Ericsson's concept enables operators to offer personalized applications to subscribers each with its own specific chain of management policies and/or services. The target architecture places an operator-controlled bridge at the customer's premises instead of a complex router, while the L3-L7 functionalities are migrated to the IP edge or into the operator cloud. Using SDN technology between the IP edge and the switch in this way offers the operator fine-grained control for dynamic configuration of the switch.

Conclusion

With its beginnings in data-center technology, SDN technology has developed to the point where it can offer significant opportunities to service providers. To maximize the potential benefits and deliver superior user experience, SDN needs to be implemented outside the sphere of the data center across the entire network. This can be achieved through enabling network programmability based on open APIs.

Service Provider SDN will help operators to scale networks and take advantage of new revenue-generating possibilities. The broader Service Provider SDN vision goes beyond leveraging split architecture to include several software components that can be combined to create a powerful end-to-end orchestration platform for WANs and distributed cloud data centers. Over time, this comprehensive software-based orchestration platform will be able to treat the overall operator network as a single programmable entity.

References

1. Ericsson, November 2012, Mobility Report, On the pulse of the Networked Society, available at: http://www.ericsson.com/ericsson-mobility-report
2. Open Networking Foundation, available at: https://www.opennetworking.org/
3. Ericsson, November 2012, Ericsson Business Review, The premium cloud: how operators can offer more, available at: http://www.ericsson.com/news/121105-ebr-the-premium-cloud-how-operators-can-offer-more_244159017_c
4. Ericsson Review, December 2012, Deploying telecom-grade products in the cloud, available at: http://www.ericsson.com/res/thecompany/docs/publications/ericsson_review/2012/er-telecom-grade-cloud.pdf
5. Ericsson Review, December 2012, Enabling the network-embedded cloud, available at: http://www.ericsson.com/res/thecompany/docs/publications/ericsson_review/2012/er-network-enabled-cloud.pdf
6. European Commission, Seventh Framework Programme, Split Architecture Carrier Grade Networks, available at: http://www.fp7-sparc.eu/
7. Elisa Bellagamba: Virtual Network System, MPLS & Ethernet World Congress, Paris, February 2012, available at: http://www.slideshare.net/EricssonSlides/e-bellagamba-mewc12-pa8

Acknowledgements

The authors would like to thank Diego Caviglia, Andreas Fasbender, Howard Green, Wassim Haddad, Alvaro de Jodra, Ignacio Más, Don McCullough and Catherine Truchan for their contributions to this article.

Elisa Bellagamba

is a portfolio strategy manager in Product Area IP & Broadband, where she has been leading SDN related activities since their very beginning. She holds an M.Sc. cum laude in computer science engineering from Pisa University, Italy.

Attila Takacs

is a research manager in the Packet Technologies Research Area of Ericsson Research. He has been the technical lead of research projects on Software Defined Networks (SDN), OpenFlow, GMPLS, Traffic Engineering, PCE, IP/MPLS, Ethernet, and OAM for transport networks. He is also an active contributor to standardization; in particular he has worked for ONF, IETF and IEEE. He holds more than 30 international patent applications; granted and in progress. He holds an M.Sc. in computer science and a post-graduate degree in banking informatics, both from the Budapest University of Technology and Economics, in Hungary. He has an MBA from the CEU Business School, in Budapest.

Joe Wilke

is head of Development Unit IP & Broadband Technology Aachen. He currently leads the SDN execution program and holds a degree in electrical engineering from the University of Aachen, Germany and a degree in engineering and business from the University of Hagen, Germany.

Telefonaktiebolaget LM Ericsson
SE-164 83 Stockholm, Sweden
Phone: + 46 10 719 0000
Fax: +46 8 522 915 99

284 23-3191 | Uen
ISSN 0014-0171