Ericsson Review

The communications technology journal since 1924 2014 • 11

Architecture evolution for automation and network programmability

November 28, 2014

ERICSSON

Architecture evolution for automation and network programmability

The target architecture of future telecom networks will be designed using sets of aggregated capabilities. Each domain will have its own set of resources that are abstracted and exposed to other domains, supporting multi-tenancy and tenant isolation. The result is a fully programmable network, that has the ability to evolve and adapt to the emerging requirements of the Networked Society.

GÖRAN RUNE, ERIK WESTERBERG, TORBJÖRN CAGENIUS, IGNACIO MAS, BALÁZS VARGA, HENRIK BASILIER, AND LARS ANGELIN

Enabled by emerging technologies like virtualization, software-defined networking (SDN) and cloud capabilities, the architecture of telecom networks is undergoing a massive transformation. This is being driven by several factors, including the need for less complex and lower-cost network operations, shorter time to customer (TTC) and time to market (TTM) for new services, and new business opportunities built on the anything as a service (XaaS) model.

The principles of the target architecture are based on separation of concerns, multi-tenancy and network programmability. As networks progress toward the target architecture supporting as-a-service models with rapid scalability capabilities and greater levels of automation, the need to focus on the basic principles will become more significant.

Full programmability of a network and its services needs to take all the building blocks of a network into consideration: how each piece will evolve; how they will interface; and how they support the structure and business processes of an operator.

SDN technologies, for example, are key enabling tools for network programmability, but to provide value they must be integrated with the end-to-end process view of the operator. Cloud orchestration technologies are also important enablers, but without proper interfaces to business management functions in place, the result would be a technically functional but commercially dysfunctional system. Well-defined technical interfaces and abstractions are critical to facilitate a split of responsibilities, support trust relationships and enable opex efficiency.

BOX A Terms and abbreviations

AAA	authentication, authorization and accounting
API	application programming interface
APN	Access Point Name
BSS	business support systems
COMPA	control, orchestration, management, policy and analytics
DC	data center
EPC	Evolved Packet Core
IGP	Internet Gateway Protocol
IPS	infrastructure and platform services
MPLS	multi-protocol label switching
MTC	machine-type communication
MVNO	mobile virtual network operator
NFV	network functions virtualization
OSS	operations support systems
opex	operational expenditure
OVF	Open Virtualization Format
PaaS	platform as a service
POD	performance-optimized data centers
R&S	routing and switching
SDN	software-defined networking
SLA	Service Level Agreement
TTM	time to market
TTC	time to customer
VM	virtual machine
vDC	virtual data center
VIM	Virtualized Infrastructure Manager

This article aims to describe the big picture of the target ecosystem, presenting an architecture description that focuses on the inter-domain interfaces, separation of concerns as well as network programmability.

The ecosystem

The target network architecture will be built using a set of critical technical interfaces that support business relations – which we call inter-domain interfaces. These interfaces mark the boundaries between the different layers or domains of a network; they support the separation of concerns, interoperability, and enable Service Level Agreements (SLAs). Administrative domains, as defined by NFV[1], are suitable for being managed as one entity from a competence and administrative responsibility point of view. As Figure 1 illustrates, there are four typical administrative domains:

- **transport;**
- **infrastructure and platform services ;**
- **access and network functions; and**
- **business and cross-domain operations.**

The target architecture – and in particular the inter-domain interfaces – serve as enablers for a multitude of domain combinations. Many other domain structures are possible, depending on the strategy and operational structure of the operator.

Administrative domains are quite physical in nature. Traditionally, they tend to consist of physical nodes with pre-integrated hardware and software functions. This, however, is changing. Together, NFV and the separation of software and hardware have brought about a new administrative domain: the infrastructure and platform services (IPS) domain. Some administrative domains – notably transport, access network and the new IPS domain – maintain responsibility for hardware and platforms, while most other network function domains – such as the Evolved Packet Core (EPC) – manage only software functions.

Even though current network architecture already includes several inter-domain interfaces, the evolution to the target architecture aims to improve multi-tenancy capabilities, as well as intra-domain and inter-domain programmability. This evolution will happen gradually and to varying degrees for each domain depending on need – in terms of value – as well as additional considerations like legacy equipment and operational processes.

Key principles of the target architecture

Developing network architecture so that it is both highly automated and programmable requires functionality to be coordinated across administrative domains. This can be achieved through a set of tools to operate each administrative domain, which have operational responsibility for the resources within the domain, as well as the ability to expose services based on these resources. In this article we refer to the combination of these operational tools as COMPA: control, orchestration, management, policies and analytics. Each term has a wider meaning than its legacy definition; all are tightly interlinked within each administrative domain, as well as having inter-domain relations. The COMPA functional groupings are illustrated in the target architecture shown in **Figure 2**.

The main principles of the target architecture are:

- separation of concerns;
- abstraction and exposure of capabilities;
- multi-tenancy;
- intra-domain programmability; and
- inter-domain programmability.

FIGURE 1 **Target architecture with example administrative domains**

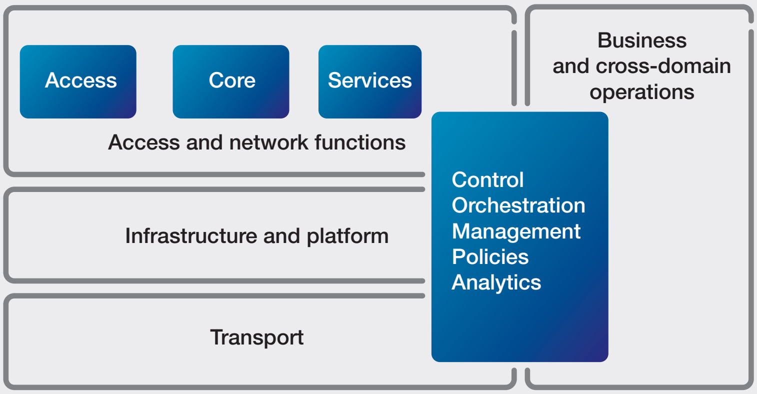

Control, orchestration and management

Management and control functions within each domain will do much the same job as they do today, but with a higher degree of automation and real-time capabilities. Orchestration enables automation across different types of resources and uses defined workflows to provide the desired network behavior – all aligned with and enabled by a policy framework that is supported by analytics insights. Creating infrastructure services is one example of where orchestration is heavily used in the IPS domain, in which processing, storage and networking resources are assigned in a coordinated manner.

Services from other domains can also be viewed as resources orchestrated in a synchronized manner with a domain's own resources to provide services in a hierarchical way. A strict framework with a common information model is required to maintain consistency across domains – illustrated by the vertical-arrow flow in Figure 2.

FIGURE 2 **Grouping of COMPA functions in the target architecture**

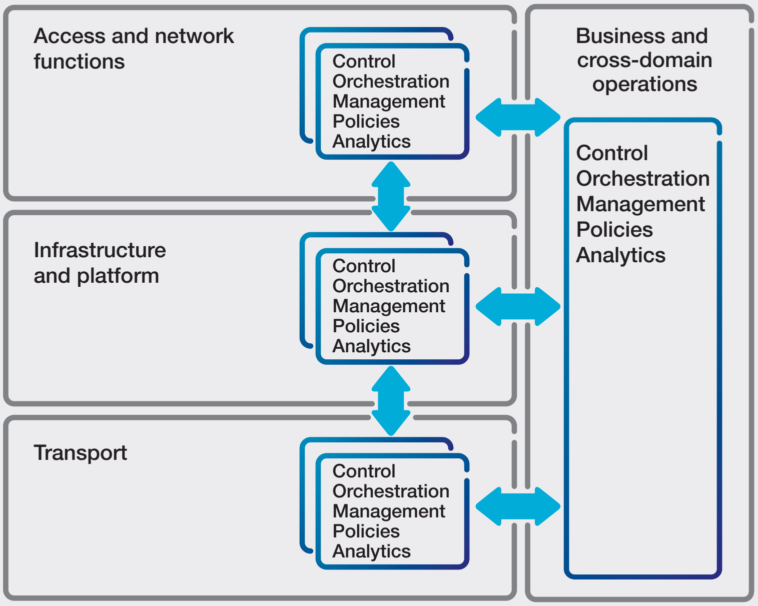

FIGURE 3 **Policy framework**

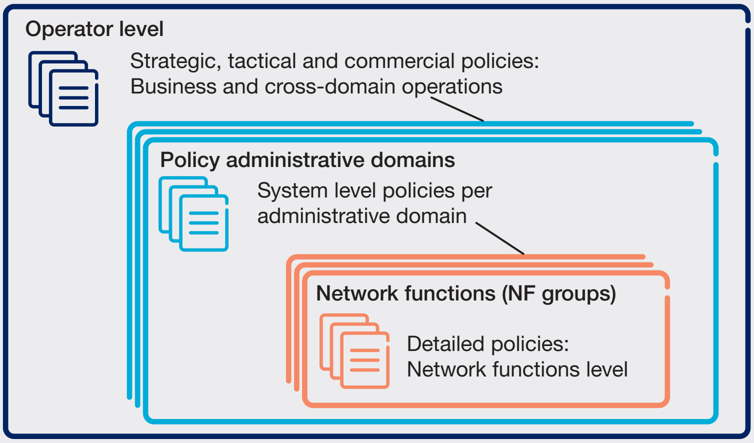
Operator level
Strategic, tactical and commercial policies:
Business and cross-domain operations
Policy administrative domains
System level policies per
administrative domain
Network functions (NF groups)
Detailed policies:
Network functions level

To offer services that draw resources from more than one domain, a cross-domain OSS/BSS function is needed. This second main flow of orchestration relates to external business offerings and how to leverage services from multiple domains. For example, an enterprise customer may require a service that combines an infrastructure service from the IPS domain with a business VPN from the transport domain – this is shown conceptually by the horizontal arrows in Figure 2.

To support service exposure, each domain needs appropriate logging tools. For example, an IPS domain will need to create and maintain data records related to usage for the infrastructure services it provides – regardless of whether it delivers these services to an external tenant or to an internal tenant (to other domains within the same operator). Many of these functions will be automated and simplified in their interfaces among staff, OSS/BSS, and resource control functions.

The policy framework

Policies are sets of rules that govern the behavior of a system. A policy comprises conditions and actions; events trigger conditions to be evaluated and actions are taken if the conditions are met. Policies are used to define a framework and set the bounds for the control-orchestration-management functions, derived from the overall business goals of the operator.

Some policies, like those that control how specific resources are used, are strictly defined and applied within an administrative domain. Other policies apply to the inter-domain interfaces, and define for example how one domain can use services from another. Such policies can be partly defined by the administrative domain delivering the service, but may also be defined by the administrative domain for business and cross-domain operations. **Figure 3** shows how policies originate from the overall business objectives of the operator and how they relate to different levels within the operator structure.

The relationship between business and network operations policies is defined by a set of meaningful operational KPIs. For example, a business policy governing the parameters of a gold subscriber service can be interpreted into specific settings for, say, QoS in the network. By factoring in the insights supplied by analytics, these operational KPIs enable a greater degree of network automation, and allow policies to govern operational decisions.

Network analytics

Analytics is therefore a key tool for increasing automation of operations. To provide insights, predictions, as well as supporting automation in other ways, analytics can be applied within an administrative domain or work in conjunction with the other COMPA functions – both in offline processing of data and for real-time stream processing. Domain competence is usually needed to understand prediction, but insights exposed from other domains or external sources could also be used as input.

Exposing analytics insights on a domain basis, and then aggregating multiple domains through a cross-domain analytics application, enables the entire network state to be analyzed; which in turn supports the definition of network-wide KPIs.

A policy engine can use network analytics to check performance-related KPIs, triggering network state updates when needed. Such requests could then be applied to the relevant network domains by the control-orchestration-management functions – possibly with some form of manual intervention.

A closed feedback loop from the control-orchestration-management functionality back to the policy engine would enable policies to learn and adapt automatically as the network environment changes.

Applying the concepts

Transport

In telco networks, the transport domain delivers connectivity services between remote sites and equipment, maintaining topology awareness and services for multiple customers – multi-tenancy. In reality, a transport network consists of a set of interworking administrative domains defined by technology, geography and ownership. The main technologies powering the delivery of connectivity services will be based on IP/MPLS, Ethernet and optical transport; in the access domain, microwave transport may also play a significant role, and IPv6 will be the dominant protocol (as IPv4 becomes more associated with legacy infrastructure). Transport network topology will become flatter with fewer packet hops, as the use of converged IP and optical transport technologies becomes more widespread[2].

Traditional connectivity services like residential broadband, mobile backhaul, and enterprise VPNs will coexist with newer services that will provide connectivity for cloud solutions, such as DC-to-DC or user-to-DC. These new generation services and the increased number of connections will drive the need for more flexible and dynamic ways to operate the transport domain.

A number of key components are needed to support evolved architectural principles and facilitate both intra-domain and inter-domain programmability. These components include SDN and network virtualization technologies[3], which allow connectivity services to be deployed and controlled in a flexible way.

Programmability in the transport domain will ensure a suitable level of resource abstraction, exposure and control so that other administrative domains can request transport services according to established SLAs. Programmability can be achieved by using northbound SDN-based interfaces, for example, and can be further increased by leveraging the benefits of data/control plane separation.

As shown in **Figure 4**, several scenarios regarding what parts of a transport node can be SDN controlled. These scenarios lead to multiple possible paths and intermediate steps to transform a traditional transport network into a network that is fully SDN-controlled – in which only a limited set of functions are local to the transport node. Using SDN controllers will not only result in the introduction of new functions and services into transport nodes, but existing control functionalities will be moved to the SDN controller – replacing current local-node implementations.

Migrating an existing transport network to an SDN-based architecture requires hybrid operational modes that apply SDN-based control capabilities onto the existing (protocol-driven local node) transport infrastructure. The capabilities that are included depend on the level of centralization versus distribution of functions that the operator chooses for its transport domain.

The resulting transport domain – in the context of packet-optical integration – combines increased programmability (enabled by SDN technologies) with the simpler, more cost-efficient IP and optical components, and is detailed in a previous Ericsson Review article[2]. The evolved transport domain enables faster service deployment and reduces operational complexity.

Infrastructure and platform services

As networks evolve, telecom solutions and systems will increasingly be built using on-demand elastic infrastructure and platform services rather than dedicated and managed infrastructure and software. To leverage the benefits of this model, a split in responsibility between the provider of such services and the users (tenants) is necessary. The provider role is taken by what we refer to in this article as the IPS domain, which is a new domain type that provides infrastructure and platform services using owned or leased resources.

One of the key services offered by the IPS domain is a structured collection of virtual computational processing, storage and networking resources, within what is referred to as virtual data center (vDC). The vDC interface separates logical telecom nodes from the actual physical infrastructure, using concepts like virtual machines, virtual network overlays, baremetal, and storage services.

Networking capabilities exposed to tenants will be rich enough to support a wide set of telco functions, including L2 and L3 VPN interworking and SDN-controlled service chaining[4]. The IPS domain can also take the administrative responsibility for common network functions (such as DNS, firewalling, DHCP, and load balancing) and offer these as services, orderable as products deployable in a vDC.

In addition, the IPS domain can also supply services to applications, providing an execution framework (PaaS) and network APIs that expose underlying network capabilities. For example, common network functions can be exposed and made programmable by applications. Inter-domain programmability and abstraction increases application development productivity and reduces lead times. In addition, the IPS domain will support migration by providing interconnectivity with non-virtualized networks as well as mixed

»

FIGURE 4 **Scenarios for control plane and data plane separation for packet, and IP/optical transport networks**

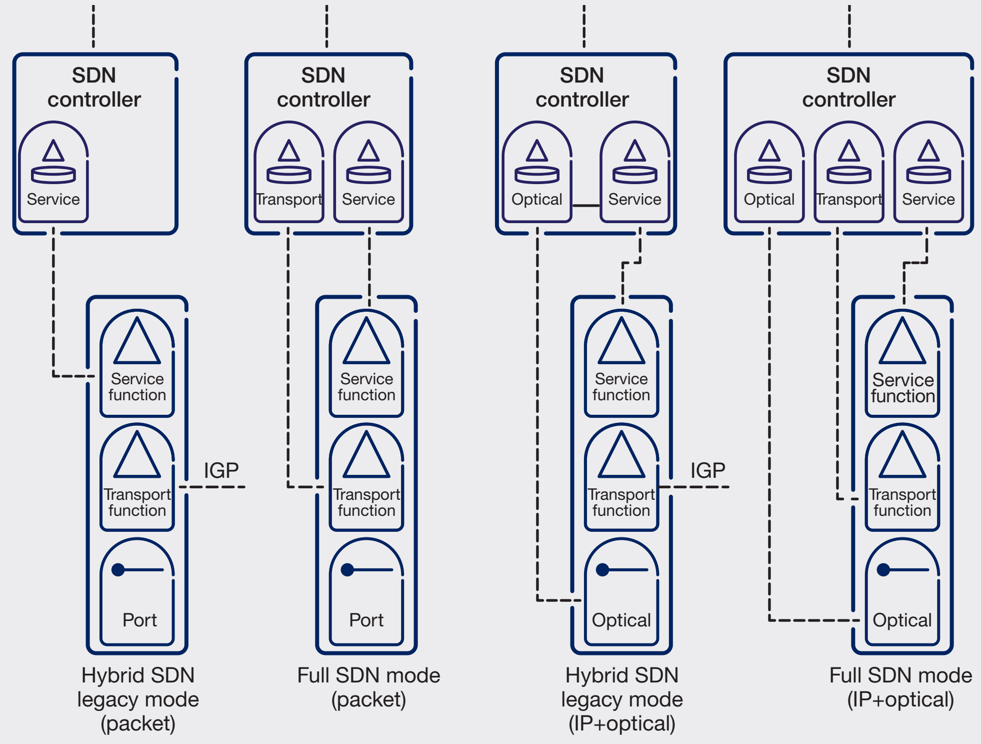

FIGURE 5 **Infrastructure and platform services domain**

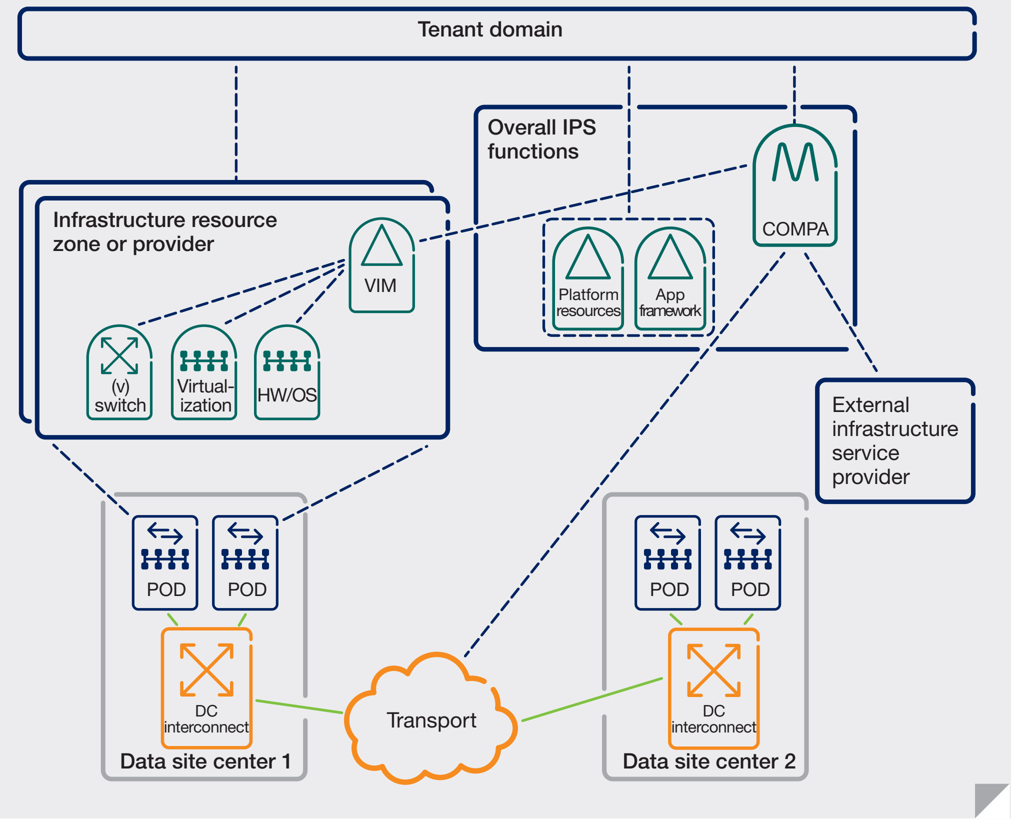

deployments of non-virtualized, virtualized and PaaS-based applications.

All the capabilities of the vDC and application services are orderable by tenants through policy-controlled inter-domain interfaces, and all of the capabilities can be requested, monitored and maintained/scaled through these interfaces. The interfaces will rely heavily on modeling of the (sometimes complex) sets of capabilities, using OVF descriptors, for example, and forwarding descriptors for service chaining.

Within the IPS domain, overall functions in the COMPA category will act across a wide set of resources in the underlying infrastructure.

Using orchestration technologies, for example, suitable abstractions can be provided to tenants using a heterogeneous set of resources – which allows tenants to manage and program resources without requiring any lower level implementation details. Policies and analytics may then be used to ensure that resources are used efficiently, while respecting SLAs and business requirements.

The physical resources that expose virtual resources to tenants may be organized into infrastructure resource zones, each with their own functions (VIM in ETSI NFV terminology) acting within the zone – such as OpenStack and SDN controllers. Some or all such zones may be external to the IPS domain. Another option is to use similar services from another IPS domain or service provider, where orchestration capabilities deliver a consolidated service. The transport domain may be used for inter-connectivity of infrastructure resource zones at different data center sites or to connect infrastructure resource zones to external networks. In both cases, the IPS domain interacts with the transport domain, based on frame agreements, to request or dynamically adapt WAN connections.

As shown in **Figure 5**, the IPS domain relies on several arbitrarily distributed DC sites, which contain a number of PODs – blocks of computational, storage and networking resources. Typically, a POD corresponds to an infrastructure resource zone. To deliver consolidated and distributed vDCs, the overall orchestrator can request resources across the PODs through their VIM functions.

The IPS domain offers abstracted services (the vDCs and application services), multi-tenancy with isolation of resources, security and SLAs associated with these services. It allows for intra-domain programmability and automation via the VIM (OpenStack), SDN for the connectivity resources and the COMPA functions for resource and service orchestration across infrastructure resource zones and to external providers. It also offers inter-domain programmability where tenants have access to interfaces for controlling – within frame agreements – their instances of the vDC and application services, supporting for example scaling, tenant SDN control or access to telco network capabilities. The interface between the IPS domain and its tenants needs to be open and, where applicable, standardized to support a full business ecosystem between IPS-domain service providers and its tenants, with a minimum amount of system integration between the two. Indeed, this appears to be one of the main tasks of the NFV forum.

Network functions

Most network functions of the logical telecom architecture shown in **Figure 6** benefit from using services from the IPS domain. The separation of network functions from platforms can result in significant operational gain – primarily through automated routines for backup and restore, capacity planning, hardware handling and a general reduction in the number of platforms to be managed. This has a direct impact on TTM for new services, which can be reduced from up to a year down to a few months as the introduction process no longer depends on platform introduction. Auto scaling of the infrastructure and platform services and programmability of the network functions removes much of the manual work associated with fulfillment, which greatly reduces the TTC.

The original design of mobile network architecture in 3GPP supports a certain level of programmability, abstraction and multi-tenancy. Standardized interfaces between the RAN, EPC and IMS domains support automation in bearer service handling and a set of MVNO solutions at various levels. The Rx interface enables rudimentary inter-domain programming to the PCRF from outside the EPC domain, while the APN structure provides a foundation for multi-tenancy. However this is not sufficient, network functions architecture is evolving to increase support for COMPA functions. Introducing the infrastructure and platform services are a significant step in this direction, but additional architectural changes and interface improvements are also part of the wider picture.

Separating network functions from the platforms allows the capacity of a given network function system – such as an EPC system – to scale up or down by simply adjusting the capacity of the vDC to achieve the wanted capacity of the EPC system. The multi-tenancy of the vDC service also means that multiple EPC systems can be instantiated in parallel in separate vDCs.

Figure 7 illustrates how deploying a multitude of EPCs in different vDCs provides full isolation of the EPC instances, inherited from the tenant isolation built into the vDC service from the IPS domain. Isolation makes both service exposure and inter-domain programmability to EPC instances safer – opening up programmability to one instance does not impact others, and exposure of data from the EPC system to a customer or partner is limited to that of the associated EPC system instance. Implementing isolation in this way minimizes risk and reduces the cost for troubleshooting faulty services.

For operations in multiple markets, one EPC system can be instantiated per market, with a central responsibility for the EPC domain, but with selected programmability suitable for the demands of the given market. This is a cost-efficient approach with consolidated competence and responsibility, while still allowing different operational entities to control selected features of the EPC system – such as rules for charging or subscription.

Instantiating a VoLTE system[5], for example, can enable an operator to offer communication services to enterprises, emergency services or any other industry with full isolation and varying degrees of programmability. To support this use case, network architecture needs to evolve to the target architecture. In particular, additional inter-domain interfaces (to enable programmability and automated orchestration) are needed to instantiate the relevant subsystems and combine them into service solutions.

The evolution of the network functions integrates well with 5G radio evolution[6]. Next generation networks will support legacy services as well as new services like enhanced mobile broadband, massive machine-type communication (MTC), as well as mission-critical MTC. Future networks will need to support a vast number and a much more diverse set of use cases. Consequently, service creation that is platform-independent and flexible, based on programmability and automation is key. A massive range of industries will depend on 5G networks – all with different requirements for characteristics, security, analytics and cost. Meeting all of these needs is a strong driver for multi-tenancy, isolation, and instantiation of services and resources.

Extending instantiation capabilities to work across multiple domains may enable novel business offerings to be created. If, for example, an instance of an EPC system is integrated with a VoLTE system instance, the two are then connected to an IP VPN, and finally

FIGURE 6 Logical telecom architecture

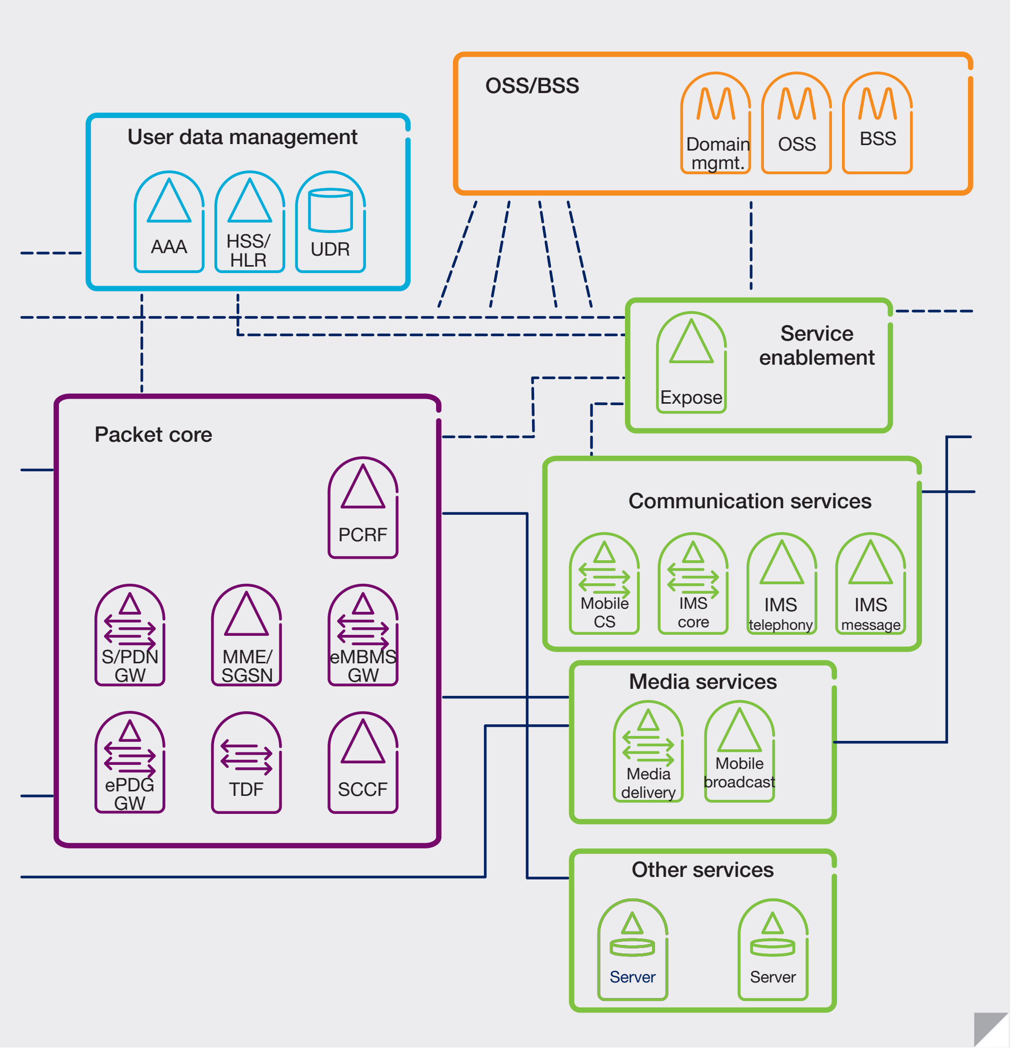

all three are associated with an isolated and SLA controlled radio-access service; the result is an isolated, and SLA-controlled logical instance of the complete network. Such logical network instances can be offered to an industry, to an MVNO or an enterprise. As each network instance is isolated, it is safe to open up interfaces to each instance to enable each customer or partner to program selected properties of the logical network instance, and to do this in real time.

To reach the point where a network can be offered as a programmable service requires a cost-efficient way to connect services – and eventually resources – from the various domains into logical network instances. As described at the beginning of this article, to connect services in such a cost-efficient way requires inter-domain programmability and more generally a network-wide architecture for cross-domain orchestration and management, while maintaining per-domain responsibility and accountability.

Conclusions

Increased levels of automation and programmability are transforming network architecture. This transformation is being driven by expected gains in operational efficiency and reduced TTM for new services, reduced TTC, and new business, as well as by the fact that enabling technologies such as virtualization and SDN are gaining maturity.

The target architecture is built on interfaces that support the principles of service and resource abstraction, multi-tenancy and programmability. Inter-domain interfaces also support business relations, as they include security and SLAs, as well as separation of responsibility and accountability.

As a first important transformation step toward the target architecture, many network functions will be managed in similar way as any other virtualized software: following virtualization management principles in line with ETSI NFV specifications. Initially virtualized network functions will be operated in parallel with legacy nodes, and DC operations as well as maintenance will be automated to a much larger degree than it is today.

In the longer term, the architecture should be able to provide the desired level of automation and network programmability. Full programmability of the network and its services requires the inter-domain interfaces as well as the domains to evolve. To achieve the full gain of the network architecture transformation, the related internal operator processes (like workflow, operation, and maintenance processes) will need to be adjusted. Technologies like SDN and cloud orchestration are crucial enablers and tools for automation

FIGURE 7 Architecture evolution

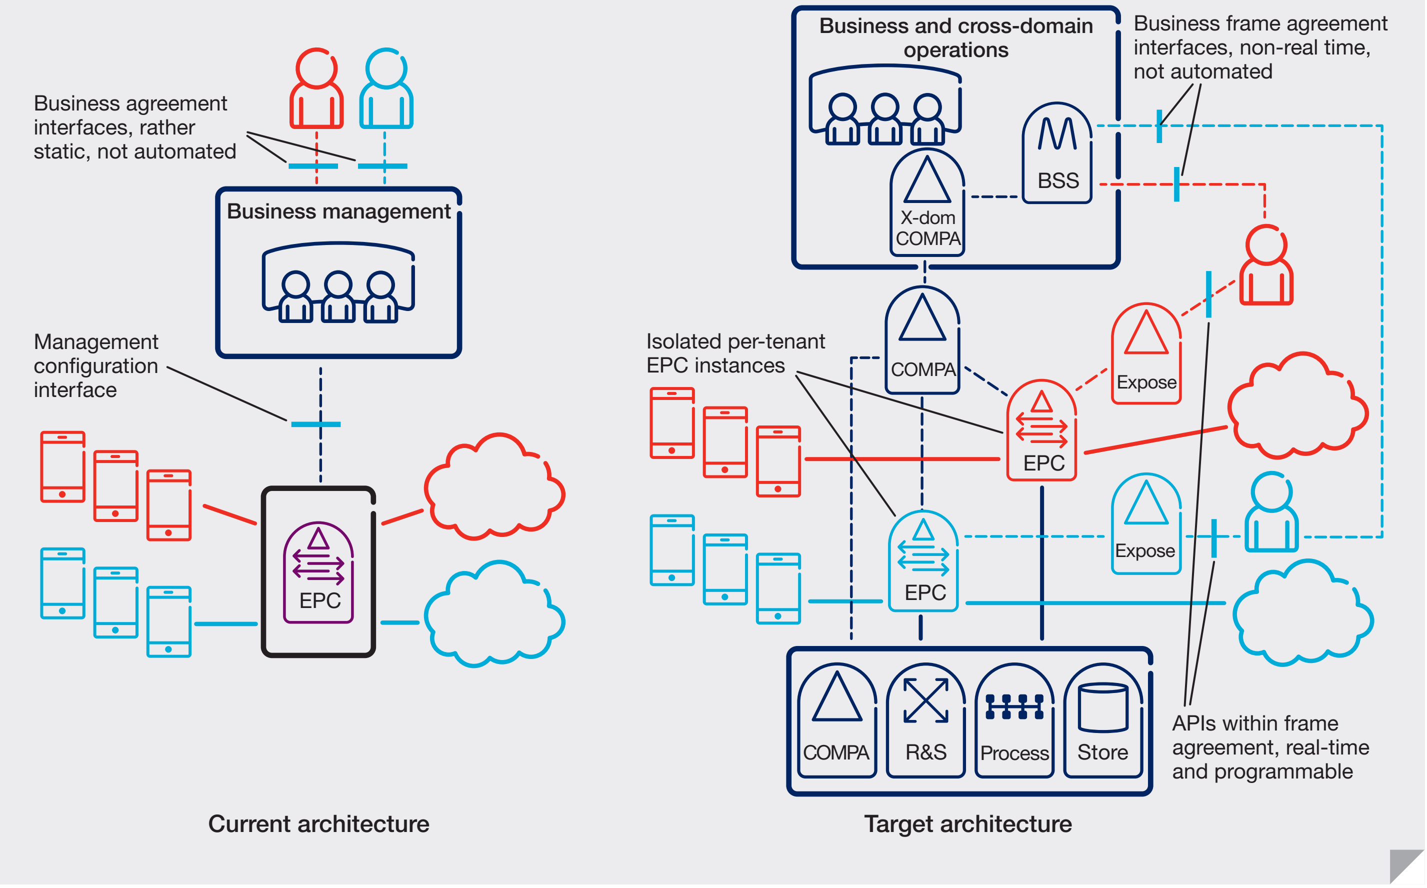
Business agreement interfaces, rather static, not automated
Business management
Management configuration interface
EPC
Current architecture
Business and cross-domain operations
BSS
X-dom COMPA
Business frame agreement interfaces, non-real time, not automated
Isolated per-tenant EPC instances
COMPA
EPC
Expose
EPC
Expose
COMPA
R&S
Process
Store
APIs within frame agreement, real-time and programmable
Target architecture

network programmability, but network operations and services also need to be controlled through operational policies linked to business policies.

Due to the impact on operator processes and potentially even the business ecosystem it is likely that the transformation will take place in a stepwise manner over a significant period of time – with different parts of the network evolving at different rates. In addition, the resulting network architecture will support 5G radio evolution and the associated use cases and requirements.❖

BOX B Main principles of the target architecture

Separation of concerns
Each domain has full responsibility over the resources and operations performed inside the domain.

Exposure and abstraction of capabilities
The abstraction of functions into APIs that are exposed as services supports domain inter-operability, which enables automation and programmability.

Multi-tenancy
Each domain offers full isolation of how the different users (tenants) use domain resources.

Intra-domain programmability
This is achieved by leveraging automation and programmability within an administrative domain through its COMPA functions.

Inter-domain programmability
Each domain exposes capabilities and services using well-defined APIs to achieve an end-to-end service offering, orchestrated by the cross-domain COMPA functionality.

References

1. ETSI, 2014, Draft Group Specification, Security and Trust Guidance, NFV ISG Spec, available at: http://docbox.etsi.org/isg/nfv/open/Latest_Drafts/nfv-sec003v111 security and trust guidance.pdf
2. Ericsson Review, May 2014, IP-optical convergence: a complete solution,available at: http://www.ericsson.com/news/140528-er-ip-optical-convergence_244099437_c
3. Ericsson Review, February 2013, Software-defined networking: the service provider perspective, available at: http://www.ericsson.com/news/130221-software-defined-networking-the-service-provider-perspective_244129229_c
4. Ericsson Review, March 2014, Virtualizing network services – the telecom cloud, available at: http://www.ericsson.com/news/140328-virtualizing-network-services-the-telecom-cloud_244099438_c
5. Ericsson Review, July 2014, Communications as a cloud service: a new take on telecoms, available at: http://www.ericsson.com/news/140722-communications-as-a-cloud-service-a-new-take-on-telecoms_244099436_c
6. Ericsson Review, June 2014, 5G Radio Access, available at: http://www.ericsson.com/news/140618-5g-radio-access_244099437_c

Torbjörn Cagenius

is an expert in distributed network architecture at Business Unit Cloud and IP. He joined Ericsson in 1990 and has worked in a variety of technology areas such as FTTH, main-remote RBS, FMC, IPTV, network architecture evolution, SDN and NFV. In his current role, he focuses on cloud impact on network architecture evolution. He holds an M.Sc. from KTH Royal Institute of Technology, Stockholm, Sweden.

Erik Westerberg

joined Ericsson from MIT, Massachusetts, the US, in 1996 and currently holds the senior expert position in system and network architecture. In his first 10 years at Ericsson, he worked with the development of mobile broadband systems before broadening his scope to include the full network architecture, serving as chief network architect until 2014. He holds a Ph.D. in quantum physics from Stockholm University, Sweden.

Henrik Basilier

is an expert at Business Unit Cloud and IP. He has worked for Ericsson since 1991 in a wide range of areas and roles. He is currently engaged in internal R&D studies and customer cooperation in the areas of cloud, virtualization and SDN. He holds an M.Sc. in computer science and technology from the Institute of Technology at Linköping University, Sweden.

Göran Rune

is a principal researcher at Ericsson Research. His current focus is the functional and deployment architecture of future networks, primarily 5G. Before joining Ericsson Research, he held a position as an expert in mobile systems architecture at Business Unit Networks focusing on the end-to-end aspects of LTE/EPC, as well as various systems and network architecture topics. He joined Ericsson in 1989 and has held various systems management positions, working on most digital cellular standards, including GSM, PDC, WCDMA, HSPA, and LTE. From 1996 to 1999, he was a product manager at Ericsson in Japan, first for PDC and later for WCDMA. He was a key member of the ETSI SMG2 UTRAN Architecture Expert group and later 3GPP TSG RAN WG3 from 1998 to 2001, standardizing the WCDMA RAN architecture. He studied at the Institute of Technology at Linköping University, Sweden, where he received an M. Sc. in applied physics and electrical engineering and a Lic. Eng. in solid state physics.

Lars Angelin

is an expert in the multimedia management technology area at Business Unit Support Solutions. He has more than 28 years of experience in the areas of concept development, architecture and strategies within telecom and education. He joined Ericsson in 1996 as a research engineer, and in 2003 he moved to the position of concept developer for telco-near applications, initiating and driving activities mostly related to M2M and OSS/BSS. He holds an M.Sc. in engineering physics and a Tech. Licentiate in tele-traffic theory from Lund Institute.

Ignacio Mas

is a system architect at Group Function Technology and an expert in network architecture. He holds a Ph.D. in telecommunications from KTH Royal Institute of Technology, Stockholm, and an M.Sc. from both KTH and the Technical University of Madrid (UPM). He joined Ericsson in 2005 and has worked in IETF standardization, IPTV and messaging architectures, as well as media-related activities for Ericsson Research. He is a member of the Ericsson System Architect Program (ESAP) and has research interests in QoS, multimedia transport, signaling and network security, IPTV and, most recently in cloud computing.

Balázs Varga

joined Ericsson in 2010 and he is an expert in multiservice networks at Ericsson Research. His focus is on packet evolution studies to integrate IP, Ethernet and MPLS technologies for converged mobile and fixed network architectures. Prior to Ericsson, he worked for Magyar Telekom on the enhancement of broadband services portfolio and introduction of new broadband technologies. He has many years of experience in fixed and mobile telecommunication and also represents Ericsson in standardization. He holds a Ph.D. in telecommunication from the Budapest University of Technology and Economics, Hungary.

Acknowledgements

The authors gratefully acknowledge their colleagues who have contributed to this article: Jaume Rius i Riu and Ulf Olsson.

Ericsson
Review

To bring you the best of Ericsson's research world, our employees have been writing articles for Ericsson Review – our communications technology journal – since 1924. Today, Ericsson Review articles have a two- to five-year perspective and our objective is to provide you with up-to-date insights on how things are shaping up for the Networked Society.

Address:
Ericsson
SE-164 83 Stockholm, Sweden
Phone: +46 8 719 00 00

Publishing:
Additional Ericsson Review material and articles are published on: www.ericsson.com/review. Use the RSS feed to stay informed of the latest updates.

Ericsson Technology Insights
All Ericsson Review articles are available on the Ericsson Technology Insights app available for Android and iOS devices. The link for your device is on the Ericsson Review website:www.ericsson.com/review. If you are viewing this digitally, you can:
download from Google Play or
download from the App Store

Publisher: Ulf Ewaldsson

Editor:
Deirdre P. Doyle
deirdre.doyle@jgcommunication.se

Subeditors:
Paul Eade and Ian Nicholson

Art director and layout:
Carola Pilarz

Illustrations:
Claes-Göran Andersson

ISSN: 0014-0171

Volume: 91, 2014

Ericsson
SE-164 83 Stockholm, Sweden
Phone: + 46 10 719 0000

ISSN 0014-0171
284 23-3235 | Uen