CHARTING THE FUTURE OF INNOVATION VOLUME 93 | #4·2016

ERICSSON TECHNOLOGY Review

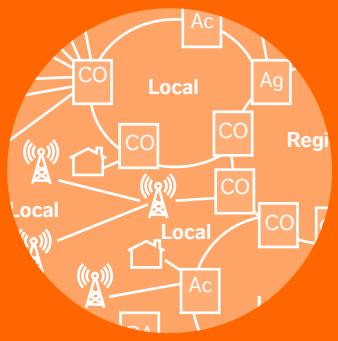

THE CENTRAL OFFICE OF THE ICT ERA: AGILE, SMART, AND AUTONOMOUS

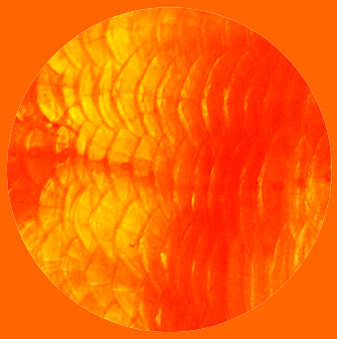

ERICSSON

THE CENTRAL OFFICE OF THE ICT ERA:: AGILE, SMART, AND autonomous

Network architecture is undergoing a massive transformation, which in turn is having an impact on the role of the central office. Enabled primarily by virtualization and SDN technologies, network architectures are becoming more flexible, with improved programmability and a greater degree of automated behavior. In combination with technology enablers such as the increased reach offered by fiber, automation of provisioning and orchestration, and improvements in the performance of generic hardware, network transformation has provided operators with the opportunity to rationalize and consolidate infrastructure. The next generation central office will introduce intelligence and service agility into the network through disaggregation.

NAIL KAVAK
ANDREW WILKINSON
JOHN LARKINS
SUNIL PATIL
BOB FRAZIER

THE CENTRAL OFFICES (COs) of fixed networks and the mobile switch offices (MSOs) of mobile operators house the networking functionality, management, and compute power needed to provide voice and data services to enterprise and residential subscribers. To route traffic efficiently, COs are distributed throughout the entire geographic region served by the network, and provide operators with a key asset: local proximity to their subscribers.

Traditionally, the location of a fixed-line CO has been determined by the reach constraints of the access technologies used in the last mile – from the CO to the subscriber (residential or enterprise). Until recently, copper was the predominant media, and so the location of the CO has been dictated by the maximum reach of the copper pairs supporting POTS equipment in the home or at the enterprise premises. Although copper is no longer the primary choice for access media (or even present in many cases), the location of COs still reflects the original distance constraints. As a result, even mid-sized cities with around a million subscribers are served by hundreds of COs, and it is still common for these to be placed in a grid-like manner, spaced a couple of kilometers apart. In rural areas with low population density, fixed-access technology reach is also the main factor for determining location, explaining why the ratio of subscribers to COs in rural areas tends to be low.

The first wave of CO consolidation and centralization came about during the digitization of POTS. Digitization resulted in a reduction in size or functionality of many city and rural COs, and in many places, they were replaced with small concentrators connected to a smaller number of more centralized COs.

The location of a CO is significant; when positioned in close proximity to users, certain services can be provided to local groups of subscribers in a highly efficient manner. This capability is one of the primary differentiating assets of the access operator.

Likewise for mobile networks, the optimal placement of an MSO takes location constraints into consideration, and the critical factor for mobile is access to the base transceiver station (BTS). Originally, cable was used as the primary media for BTS access, shifting in recent years to use of high-capacity fiber TDM circuits. So, for the same geographic area and subscriber base, MSOs tend to be more centralized compared with fixed COs.

Today, a medium-sized city can be served by just one or two MSOs, but possibly hundreds of COs. In rural areas, however, MSOs tend to be sparse or even nonexistent. Operators running converged fixed and mobile networks tend to house MSOs within existing COs, rarely opting for new builds in dedicated locations.

“THE LOCATION OF A CO IS SIGNIFICANT; WHEN POSITIONED IN CLOSE PROXIMITY TO USERS, CERTAIN SERVICES CAN BE PROVIDED TO LOCAL GROUPS OF SUBSCRIBERS IN A HIGHLY EFFICIENT MANNER.”

Figure 1 illustrates the local, regional, and nationally tiered structure of COs. Fixed COs have two or more progressively centralized tiers, which originally provided inter-office calling capability to avoid the need for a full CO mesh. Higher-tier COs have extensive transmission trunking from lower-tier and access COs, which is significant, as this architecture may be utilized for the placement of next generation central offices. MSOs can be colocated with a subset of COs or be deployed independently as local and regional COs. End sites

Terms and abbreviations

ACL — access control list | **API** — application programming interface | **ARPU** — average revenue per user | **BGP** — Border Gateway Protocol | **BNG** — Broadband Network Gateway | **BSC** — base station controller | **BSS** — business support systems | **BTS** — base transceiver station | **CDN** — content delivery network | **Clos** – Non-blocking, multistage switch fabric formalized by Charles Clos | **CMS** – cloud management system | **CO** — central office | **COTS** — commercial off-the-shelf | **CPU** — central processing unit | **DOCSIS** — Data Over Cable Service Interface Specification | **DSL** — digital subscriber line | **GPON** — gigabit passive optical network | **HLR** — home location register | **I/O** — input/output | **IGP** — Interior Gateway Protocol | **IoT** — Internet of Things | **ISP** — internet service provider | **M2M** — machine-to-machine | **MAC** — media/medium access control | **MME** — Mobility Management Entity | **MPLS** — multi-protocol label switching | **MSO** — mobile switch office | **NETCONF** — protocol to install, manipulate, and delete the configuration of network devices | **NFV** — Network Functions Virtualization | **NGCO** — next generation central office | **NIC** — network interface card | **NMS** — network management system | **NVGRE** — Network Virtualization using Generic Routing Encapsulation | **ODL** — OpenDaylight | **OLT** — Optical Line Termination | **ONIE** — Open Network Install Environment | **OSS** — operations support systems | **OTN** — optical transport network | **OTT** — over-the-top | **PON** — passive optical network | **POTS** — plain old telephone service | **P-GW** — packet data network gateway | **P/S-GW** — packet data network/serving gateway | **RNC** — radio network controller | **SDN** — software-defined networking | **SFP** — small form-factor pluggable | **S/GGSN** — serving/gateway GPRS support node | **vBNG** — virtual Broadband Network Gateway | **VIM** — virtual infrastructure manager | **VoD** — video on demand | vSwitchvirtual switch | **VXLAN** — Virtual Extensible LAN | **XaaS** — Sanything as a service | **XMPP** — Extensible Messaging and Presence Protocol

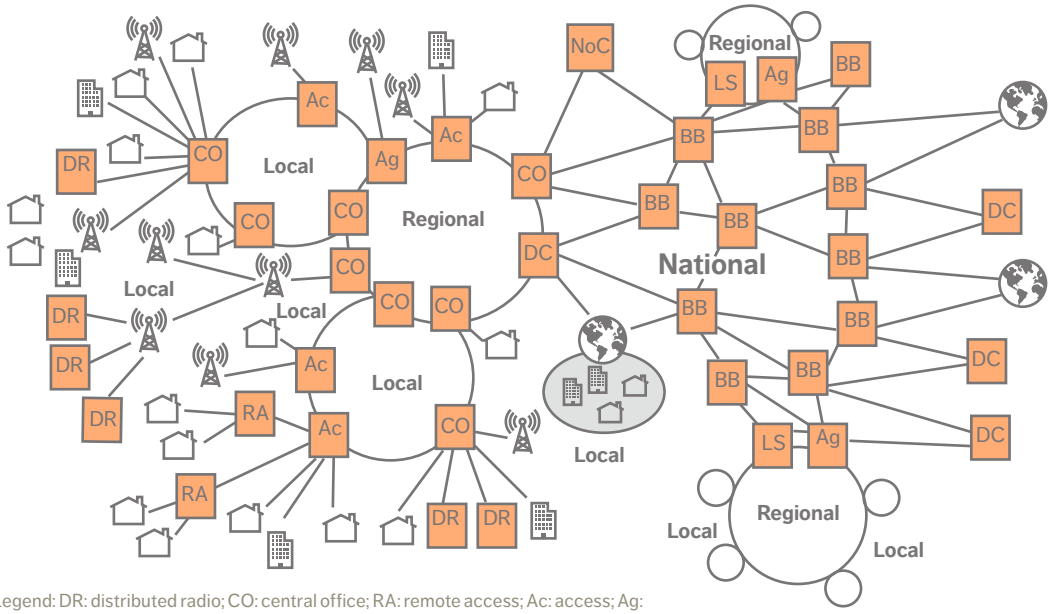

Legend: DR: distributed radio; CO: central office; RA: remote access; Ac: access; Ag: aggregation; DC: data center BB: backbone; LS: local switching

Figure 1: CO tiers and distribution

connect to COs and MSOs through intermediate transport aggregation sites.

The COs house POTs and DSL/PON access equipment, and to a lesser extent, IP/Ethernet routing and switching capabilities for residential and enterprise services. MSOs house radio aggregation nodes, such as BSCs and RNCs, as well as transport switches. Some of the MSOs may house additional 3GPP core functions such as S/GGSN and P/S-GW, as well as control-plane functions such as MME and HLR serving multiple geographic areas. Other core functions can, however, be placed elsewhere, for example in purpose-built regional or national DCs.

In this second wave of CO consolidation, the fundamental internal structure and functionality provided at each site will change, and use of new technologies will either result in fewer sites or greater capacity. The term next generation central office (NGCO) has been adopted by the telecom industry to refer to the future central offices that will support both fixed and mobile operations. Compared with its current CO counterpart, the NGCO will be able to serve more subscribers, implement access functions in a more IT-centric way, and support and locally house new, flexible data services. The NGCO will function like a highly automated mini data center, requiring less space, power, and cooling than the set of traditional COs it replaces.

Why transform?

In addition to the constant need to reduce opex and capex, fixed and mobile operators continually face new challenges as technology and user demands change. Network transformation and changing

subscriber traffic patterns have created new challenges in terms of the services operators offer, and perhaps more significantly, the services that operators would like to offer, and how to provide them in the shift toward the more attractive anything-as-a-service (XaaS) business model.

The shift from voice to data services and the corresponding massive increase in OTT traffic have put pressure on networks. Changes in user behavior, with preferences shifting to use of bandwidth-hungry data services, and video consumption require a revolutionary change in the way existing CO- and MSO-based network architectures are structured.

Traffic patterns and demands

The annual growth rate of traffic carried by mobile and fixed networks has risen massively over the past five years. In addition to increasing traffic volumes, meeting the ever more stringent demands placed on network performance characteristics, such as latency, is necessary to support emerging industry applications. Technology improvements made in fixed-network access and the mobile industry (as 5G systems evolve) will enable networks to cope with growing traffic volumes and performance demands. But, as network capabilities increase, user expectations and the demand for more capacity and bandwidth will also inevitably rise.

The increase in traffic volumes and performance demands can be predicted and planned for, but changing traffic patterns due to changing subscriber habits is complicating network architecture in a new way. As networks become more flexible, user-to-user and machine-to-machine flows will become more widespread, adding new dimensions to the traditional user-to-server traffic-flow pattern. Factor in the massive expansion of the Internet of Things (IoT) and the result will be an explosion in the number of flows and routes that networks will need to support.

With static or declining ARPU, the question facing many operators is how to invest in networks so they meet constantly rising performance demands.

Technology provides some useful steps that can help answer this question. For example, where possible and necessary to meet latency requirements or lower backhaul costs, self-served and partner content, such as video, and subscriber-associated IP service delivery points – P-GWs, BNGs, and multi-service edge routers – can be moved closer to the user. Traffic not served by the access operator can be offloaded to other ISPs, transit carriers, or OTT content providers that are closer to the access domain, rather than hauling it back to more centralized interconnection points. Similarly, instead of hubbing enterprise transport traffic through large centralized routing points, a more optimal way to route this type of traffic is through distributed routing points in the network. Shifting traffic around like this will dramatically alter the ratio of locally terminated traffic to transit traffic and requires the NGCO to provide support for routing and service functionality well beyond the capabilities of the traditional CO.

Efficient rollout of services

To take advantage of the revenue streams created by massive traffic volumes, tough performance targets and new traffic patterns, networks need to be able to support efficient rollout of services. Network flexibility is key here, enabling operators – and indirectly subscribers – to modify services to match their evolving needs, scale them easily, and be able to specify and change the location of service instantiation. Provisioning mechanisms need to be highly efficient, low opex and capex are essential, and, as time to market is crucial, high feature velocity is vital.

Access operators offer end services such as web applications, CDNs with their associated content caches, and bump-in-the-wire services including parental control filtering, as well as transport services such as enterprise connectivity or internet access, or a combination of both. More advanced services require support for service chaining that can be dynamically customized on a per-subscriber basis.

Public and private cloud-based XaaS is an attractive offering for both enterprise and non-enterprise customers, but requires support for multi-tenancy environments.

> THE INCREASED PENETRATION OF FIBER IN THE LAST MILE IS PERHAPS THE MOST SIGNIFICANT FACTOR IN THE SHIFT TOWARD FEWER AND MORE CENTRALIZED NGCOS

By appropriately locating these services in NGCOs, carrier networks will become part of a distributed and intelligent cloud resource, supplementing larger, centralized data centers.

Software development and deployment life cycle

A typical service life cycle starts with development and verification before moving on to wide-scale deployment in the network.

Efficient service life cycle depends on two key factors: short time to market and deployment flexibility. Time to market can be minimized through a homogeneous software environment that enables deployment on existing network infrastructure without the need for hardware modification. Deployment flexibility is needed to enable elastic capacity scaling, dynamic service chaining, and the deployment of services in new locations.

Key to implementing these factors in the NGCO is virtualization of the compute platform on which services run, so that the traditional coupling of software to specific hardware can be removed. Decoupling provides a homogeneous development and deployment environment that is suited to an automated life cycle.

Technological enablers

Fiber reach

The increased penetration of fiber in the last mile is perhaps the most significant factor in the shift toward fewer and more centralized NGCOs. Connectivity over the last mile may be delivered by a PON. This might come in the form of fiber, or as a hybrid solution in which a relatively short copper extension using VDSL or DOCSIS technology extends the fiber from the NGCO to the curb.

As an enabler, fiber applies primarily to the central offices for fixed services, as mobile offices already tend to be positioned to operate with long-reach access technologies.

Virtualization

As it decouples applications from the underlying hardware platform, virtualization is one of the key enablers for flexible service and function deployment.

With good orchestration, virtualization technologies enable most types of workloads to be consolidated on common multi-core compute platforms. Further reduction of hardware in the NGCO can be achieved by pooling workloads on a common compute resource, and additional power savings can be gained through dynamic workload reassignment.

The significance of virtualization in future carrier networks is clearly reflected by the massive effort being put into this area by operators, vendors, and standardization bodies. The heightened focus on all aspects of virtualization bodes well for the acceleration of its adoption.

Automating the VNF life cycle

Automated orchestration of virtual functions' instantiation, capacity elasticity, and function termination are critical network capabilities that enable functions to be deployed quickly and flexibly in multiple, geographically distributed NGCOs.

Orchestration is central to the operation of any virtualization environment offering multi-tenancy – whether it is for an operator's many internal tenants, or external residential and enterprise tenants.

Compute performance

The continuous improvements in compute performance can be attributed to a number of different technologies. Cores, for example, have

become faster, the core per socket ratio has risen, on-chip caches have become both larger and faster, and access times to peripheral memory and storage have dropped dramatically. Today, it is fairly common for an individual CPU to contain tens of cores, each running at 3GHz on COTS hardware, with single, dual or quad sockets. In addition, I/O speeds have increased, enabling modern servers to support dual (and possibly more) 40Gbps NICs.

The increases in compute and I/O performance have in turn widened the set of functions that might benefit from virtualization. And so, network design is no longer restricted to the virtualization of traditional IT and control-plane intensive workloads, but can be expanded to include traditional telecom network functions that demand high user-plane performance, such as virtual routers and virtual subscriber gateways including virtual BNGs and P/S-GWs.

As compute capabilities continue to improve, an equivalent reduction in the hardware footprint of access functions will occur. This not only brings benefits in terms of cost and environmental impact, but also enables functions that benefit from proximity to the user, previously deployed in more spacious DCs, to be distributed and deployed in the NGCO.

DC switching fabric

To virtualize network functions and other workloads as far as possible, the NGCO obviously needs appropriate compute and storage capacity. Emerging DC fabrics – based on merchant silicon leaf-and-spine switches – that are scalable, and offer high capacity at low cost, provide just the right kind of internal network design between compute-and-storage components and the physical WAN and access gateways.

Most NGCO fabrics will be configured as non-blocking Clos [1] networks, possibly with under-subscribed dimensioning, even though such a structure is not strictly required.

Software-defined networking

Applying the concepts of SDN to a network makes it centralized, dynamically provisioned, and programmable. The agility and flexibility SDN offers will be critical in providing new and multiple-service operators with the capability to offer whatever services they like to their subscribers.

Key architectural components

Figure 2 shows the location of the NGCO and how it is connected to the fixed and mobile services it offers to subscribers through the various access domains. The diagram also includes an abstract representation of the internal structure of the NGCO and its connections deeper into the network. The orchestration component manages the functions and infrastructure of the internal office as well as certain external entities such as access routers.

Infrastructure

The NGCO infrastructure consists of three major components:

- the switching fabric that links all other components together
- gateways – to the access domain and the WAN
- servers and storage

Initially, non-virtualized bare metal appliances that perform specific functions will also be part of the infrastructure. These appliances might be incorporated into the gateways or be implemented on separate hardware platforms, depending on the capacity of the gateway and how well the hardware performs.

Switching fabric

The structure of an NGCO may use overlay/underlay design principles or adopt a more traditional approach. In an overlay/underlay design, the switching fabric forms the underlay and is agnostic of service endpoints. In traditional architectures, the switching fabric is fully aware of the service endpoints. The size and scale of the fabric varies according to the requirements and location of the office. For example, a small NGCO serving tens of thousands of users may consist of just a few switches and support a minimum set of

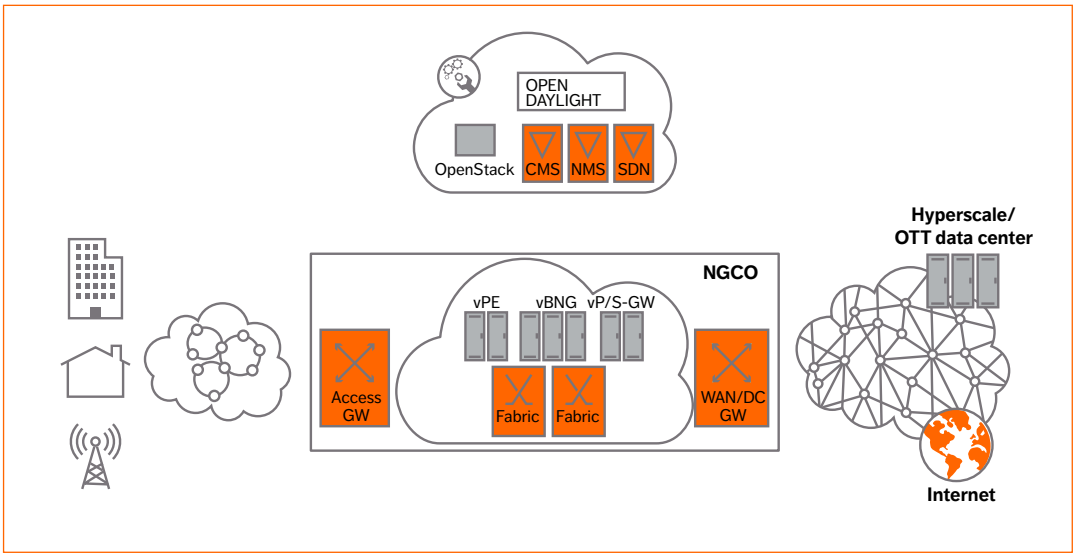

Figure 2: The NGCO in the operator's network

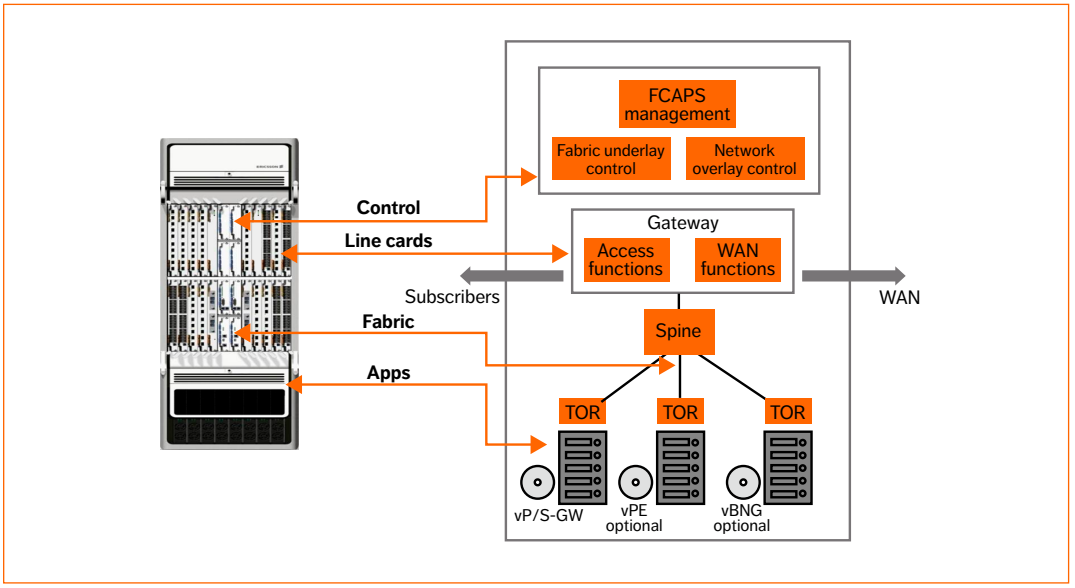

Figure 3: Disaggregation of routing functions

local functions, whereas larger offices may include a switching fabric capable of supporting extensive local services for millions of subscribers.

The structure of the fabric, especially when it comes to larger offices, is likely to be based on common data-center design practices, with an underlay Clos architecture, using a cluster of leaf-and-spine switches with same-length links, offering potentially deterministic delay and latency. In a Clos underlay, load balancing within the fabric is achieved by utilizing the multiple paths between source and destination. Either centralized SDN controllers or distributed routing protocols such as BGP or IGP will be used to build the forwarding, routing, and switching tables. To build the fabric underlay for large-scale NGCOs, the industry preference is leaning toward the use of distributed routing protocols, as they are simplistic and have a proven track record.

Merchant silicon-based white boxes can be used for fabric switches, especially when providing a simple underlay. These boxes tend to be less capable but often have a lower price-to-bandwidth ratio than traditional switches. White boxes offer entirely decoupled networking OS and hardware, and by using a tool such as the Open Network Install Environment (ONIE), for example, the installed network OS can be easily swapped out with another one – allowing operators to load the OS of their choice onto installed hardware. So, white boxes not only contribute to reducing costs; they perhaps more significantly provide network programmability and flexibility.

Shown in *Figure 3*, the NGCO fabric conceptually represents a disaggregated router that can be readily scaled out by adding leaf-and-spine switches as needed. The fabric may need to support a number of underlay technologies including IP and Ethernet, and MPLS may be required, especially in carrier domains, to ensure operational simplicity and seamless end-to-end interoperability with the installed base.

In the event of a switch failure, the fabric automatically reroutes traffic through the remaining switches until the failed switch has been manually replaced and auto-configured by a fabric manager, allowing the system to operate without having to wait for a maintenance window.

Optimum traffic management requires a holistic and real-time view of the available network bandwidth and traffic patterns. Flow statistics are collected at regular intervals, and when analyzed, provide the information needed to detect and avoid congestion, guarantee better utilization of fabric resources, and administer prioritization policies.

Gateways

Access and WAN gateways act as infrastructure gateways, and tend to be connected to special leaf nodes. The WAN gateway function could alternatively be implemented using spine switches.

Access gateways that terminate customer access links may require extended capabilities such as deep buffers, traffic management and other more advanced QoS capabilities, large forwarding tables, and ACLs that are not usually present in merchant silicon-based white boxes. Access gateways terminate different access technologies such as DOCSIS and GPON OLT. OLT functions can be virtualized with the MAC layer and the optics separated from the control- and management-plane software. The hardware part of the gateway can be implemented on a small SFP form factor, while the software part can be virtualized and hosted on any server within the CO.

Using a variety of communication protocols (such as IP, MPLS, and OTN/WDM), WAN gateways connect central offices with other NGCOs and COs, central and regional data centers, as well as other carriers and the wider internet.

Compute and storage

The geographic closeness of the NGCO to users provides a strong incentive to house certain functions and services that benefit from this proximity in the NGCO. Compute and storage resources exist in the NGCO to run virtualized network functions such as vBNG and vP/S-GW, as well as more traditional services such as VoD, with local caching.

The general purpose nature of compute

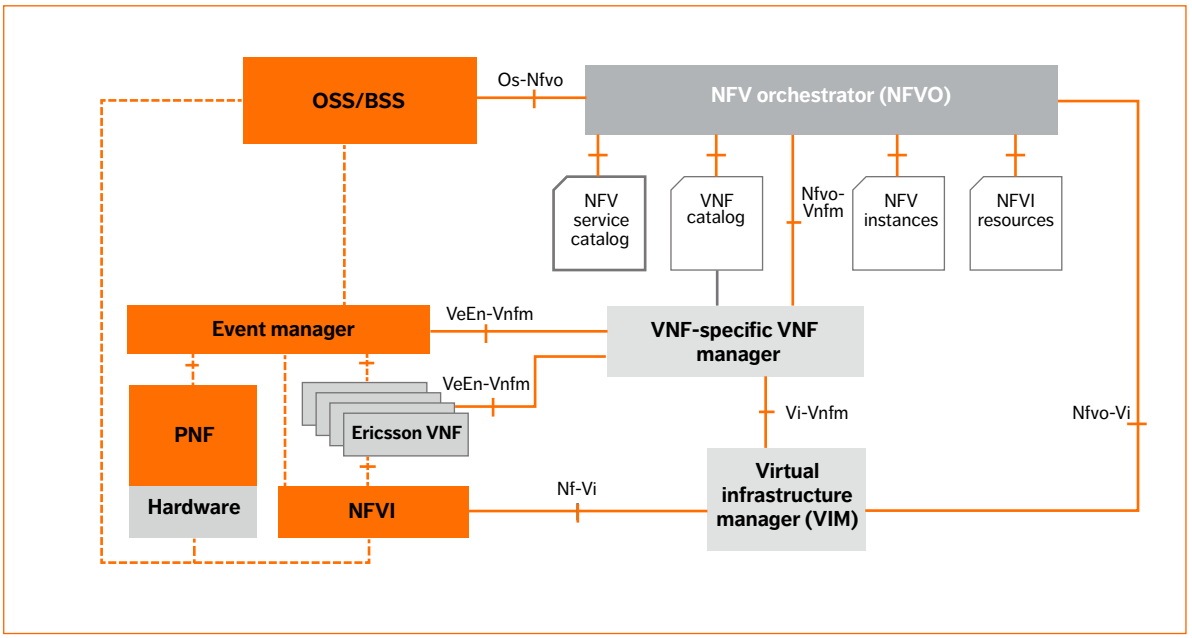

Figure 4: ETSI NFV reference architectural framework

resources deployed in the NGCO is key, as any network function or service can be instantiated on them, supporting the break away from traditional hardware and software coupling.

The amount of compute and storage located in a given NGCO will depend on its size and operator preferences for centralization versus decentralized function deployment. Offering cloud services, for example, requires additional compute and storage, which in turn increases the size of the NGCO.

Overlay services

If the NGCO implements overlay services using an underlay switching fabric, an overlay encapsulation technique is required. This technology can also be used to provide tenant isolation to operator-internal stakeholders and subscriber isolation for NGCOs with cloud services.

Common encapsulation technologies include VXLAN, NVGRE, and MPLS VPNs, and can be implemented virtually in vSwitches or in hardware on leaf and possibly gateway nodes if higher performance is required.

Regardless of the location and type of overlay technology used, configuration will be automated by an overlay controller coupled through northbound APIs to the automatic provisioning of any tenant-related functions. For example, the overlay controller could be ODL-based coupled to OpenStack through Neutron APIs. The same APIs can be used by additional applications such as the OSS/BSS. The overlay controller communicates with virtual

network switches or bare metal devices (gateways and leaf switches, for example) preferably through open southbound interfaces such as OpenFlow, XMPP, or NETCONF.

Virtualized network functions

In today's COs, traditional network functions and workloads, such as caches and webservers, run on vertically integrated platforms. In the NGCO, these elements will be run as virtualized network functions on COTS hardware.

NFV technology makes it easier to create and scale separate logical nodes and functions, and if necessary, these elements can be isolated for use by a specific tenant. This is the concept of network slicing. Network slices are individually designed to meet a specific set of performance requirements tailored to the application running on the slice. The virtual infrastructure of a slice is isolated from other slices to ensure that all slices of the network run efficiently and performance targets are met. The NFV approach provides the flexibility needed to provision network resources on demand, and to tailor slices to specific use cases, enabling operators to deliver networking as a service. The beauty of network slices lies in their ability to be optimized to suit the application. In other words, high-availability services can run on slices optimized for resilience to hardware and software failures, whereas an M2M signaling-intensive application, for example, can run on a low-latency, low-bandwidth slice.

Automation

In the NGCO, all key operational components are automated. This removes the need for manual configuration, which is prone to error, costly, and time-consuming.

The fabric manager oversees the automated parts of the NGCO, configuring and managing the underlying fabric switches, and supervising the performance of the fabric. The fabric manager continually and automatically monitors the physical fabric node-and-link topology, it validates the physical cabling, and configures leaf-and-spine switches with associated protocols and policies. The fabric manager may use DevOps tools such as Chef or Puppet for initial configuration and software management tasks (LLDP configuration, management addressing, and OS component upgrades), after which programmatic interfaces such as NETCONF/YANG can be used to configure network protocols, QoS policies, and statistics on the interfaces. For centralized SDN-based cases, the fabric manager can use OpenFlow to configure the necessary forwarding entries in the underlay switches.

> "THE GEOGRAPHIC CLOSENESS OF THE NGCO TO USERS PROVIDES A STRONG INCENTIVE TO HOUSE CERTAIN FUNCTIONS AND SERVICES THAT BENEFIT FROM THIS PROXIMITY IN THE NGCO"

Service orchestration

Service orchestration automatically instantiates applications and configures network services according to service-level specifications. Automation of these tasks can dramatically reduce the time to instantiate or add new devices or services to the network, which increases network agility, making real-time service provisioning possible.

Migration

For most operators, the migration of network architecture from the current CO deployment to one based on fewer NGCOs will be gradual. While some NGCOs will be built as greenfield deployments, for the most part, existing COs will evolve, requiring the coexistence of decoupled SDN/NFV equipment, together with traditional, tightly coupled hardware and software. During the migration/coexistence period, management and orchestration components need to be able to support the heterogeneous (coupled/decoupled) environment; by, for example,

abstracting the differences between the two architectures, and using common northbound interfaces to other systems, such as end-to-end service orchestration and OSS/BSS.

Throughout the period of coexistence, network functions will be physically and virtually instantiated with capacity and subscribers pooled across both, and as *Figure 4* shows, orchestration systems will be required to support both traditional and decoupled architectures.

Conclusions

Network architecture is undergoing a massive transformation in terms of increased levels of automation and programmability. This transformation has been enabled by a number of technologies, but primarily by the disaggregation of software and hardware. The transformation is being driven by new business opportunities, expected gains in operational efficiency, and the need for rapid time to market for services. As the underlying technologies – virtualization and SDN – become more mature, the rate of transformation will rise.

The next generation central office, or NGCO, has been designed to take advantage of the gains brought about by a decoupled network architecture. The benefits for operators come in the form of network intelligence, flexibility, and ease of scalability, all of which bring opex and capex benefits.

The NGCO is basically a mini data center that provides converged fixed and mobile services. Compared with a traditional CO, the NGCO can serve a larger subscriber base across a wider geographic area. The NGCO has been brought about through:

- reduced CO density, as a result of greater distances achievable by fixed access technologies
- the introduction of SDN/NFV technologies
- advancements in hardware technologies in terms of low-cost, high-throughput switches
- infrastructure automation and service orchestration

Architecturally, deploying the NGCO as a mini data center introduces a greater level of intelligence into the network in a distributed fashion, as applications are replicated, or shifted, from centralized data centers out to NGCOs. Compute resources in the NGCO can be used for running applications such as rich media and rendering, or latency-sensitive gaming apps. With these capabilities, the NGCO will become part of a distributed, intelligent cloud resource.

The NGCO brings with it a number of savings, requiring less space, power, and cooling than the sum of the individual traditional COs they replace. On-site staffing requirements should be reduced, as provisioning and many aspects of maintenance are controlled remotely and automated. Overall, the NGCO will result in fewer central offices or increased access coverage and service consolidation, with reduced need for new real estate as equipment continues to compact.

References:

1. **Bell Labs Technical Journal, 1953, A Study of Non-Blocking Switching Networks, Charles Clos, abstract, available at:** *http://ow.ly/YGZ2F*

THE AUTHORS

Nail Kavak

◆ joined Ericsson in 2000, and is currently working as principle architect for the system and technology group in Development Unit IP. He has in-depth experience in the design and deployment of IP/MPLS and

optical networks for carrier networks. Most recently, he has managed a number of network transformation projects for Tier 1 operators in the DC Networking space. He holds an M.Sc. in computer science and engineering from Linköping University, Sweden, and a technical licentiate from the KTH Royal Institute of Technology in Stockholm. *https://www.linkedin.com/in/nail-kavak-8ba9481*

Andrew Wilkinson

◆ is an expert in IP networking at Ericsson's Development Unit IP. He holds an M.Sc. in telecommunications from the University of London. He joined Ericsson in 2011

having previously worked for mobile network operators in Europe and North America. https://www.linkedin.com/in/andrew-wilkinson-0b377712

John Larkins

◆ is a senior director of technology at Ericsson's IP Design Unit in San Jose, California, where he is responsible for technology

evolution, including network and systems architecture solutions ranging from ASIC requirements definition to product implementation architectures and collaboration with network operators on future target network architectures. *https://www.linkedin.com/in/larkins*

Sunil Patil

◆ is a principal engineer in IP networking at Ericsson's Development Unit IP. He

joined Ericsson in 2000, where he has worked on architecture, design, and development of multiple IP routing products. His current focus is on driving technology innovation in the areas of SDN, orchestration, NGCO, and data center networking for IaaS, PaaS, and CaaS. He holds an M.Sc. in computer networking from North Carolina University, the US, and an M.B.A. from Duke University. *https://www.linkedin.com/in/sunilbpatil*

Bob Frazier

◆ is an expert in IP system architecture at Ericsson's

Business Unit Cloud & IP. He holds a Ph.D. in electrical engineering from Duke University in North Carolina, the US. He joined Ericsson in 2007 and has worked in IETF, IEEE, and Broadband Forum standardization. His current interests are IP software architecture and data center networking. *https://www.linkedin.com/in/bob-frazier-a961572*

ISSN 0014-0171
284 23-3280 | Uen

© Ericsson AB 2016
Ericsson
SE-164 83 Stockholm, Sweden
Phone: +46 10 719 0000