


# 2019 Next-Generation Central Office Report

RESEARCH BRIEF



# Table of Contents

Market Overview: The Move to Edge Computing . . . . .	1
The Edge Computing Renaissance with a Halo of Cloud Computing . . . . .	1
The Changing Carrier Infrastructure Landscape . . . . .	2
NGCO Characteristics and Technologies . . . . .	3
Hardware Components and Design . . . . .	5
NGCO Software Foundation and Attributes . . . . .	6
Benefits of NGCOs . . . . .	8
Recommendations and Next Steps for Providers and Enterprises . . . . .	9


Research Briefs are independent content created by analysts working for AvidThink LLC. These reports are made possible through the sponsorship of our commercial supporters. Sponsors do not have any editorial control over the report content, and the views represented herein are solely those of AvidThink LLC. For more information about report sponsorships, please reach out to us at [research@avidthink.com](mailto:research@avidthink.com).

**About AvidThink™**

AvidThink is a research and analysis firm focused on providing cutting edge insights into the latest in infrastructure technologies. Formerly SDxCentral’s research group, AvidThink launched as an independent company in October 2018. Over the last five years, over 110,000 copies of AvidThink’s research reports (under the SDxCentral brand) have been downloaded by 40,000 technology buyers and industry thought leaders. AvidThink’s expertise covers Edge and IoT, SD-WAN, cloud and containers, SDN, NFV, hyper-convergence and infrastructure applications for AI/ML and security. Visit AvidThink at [www.avidthink.com](http://www.avidthink.com).

# All-in-One Switch-Router on FPGA SmartNIC

- ✓ Stop Burning CPU Cores on Networking and Security!
- ✓ Reduce TCO on Edge Deployments
- ✓ Field-Proven Technology, Over 100M End Users



**ETHERNITY**  
NETWORKS

# 2019 Next-Generation Central Office Report

## Market Overview: The Move to Edge Computing

When he said, "History doesn't repeat itself, but it often rhymes," Mark Twain<sup>1</sup> obviously wasn't thinking of IT, but his quip couldn't be more apropos since the history of IT is one of cycles; none more significant than the oscillation between centralized and distributed systems. The pattern from centralization to distribution is notably resurgent in telecom infrastructure where a nexus of technological, business, and consumer changes promises to stress the demand for capacity and services beyond what traditional central office designs can handle.

The past decade or so has been marked by IT consolidation. Initially, as the server side of client-server computing did more of the work while clients became mobile portals to a set of always-available central services. Over time, those central servers evolved into complex enterprise systems that looked more and more like the mainframes they displaced. Server farms ultimately evolved into scalable cloud infrastructure and platform services that could be shared among thousands of applications and millions of users.

Cloud infrastructure continues to improve and evolve and has become the technical foundation for next-generation telecom systems and central offices. However, the downsides to massive, highly centralized infrastructure are already being felt by users and service providers, as an explosion in streaming content, mobile users and connected devices has exacerbated the performance limitations of long-haul networks. First to feel the scalability limits of service centralization and wide area networks (WAN) are the service providers (SPs): telecommunications providers, cable multiservice operators (MSOs), and internet service providers (ISPs) responsible for delivering content and applications. Sadly, strains on their existing infrastructure are but a prelude to more severe ramifications once 5G wireless technology opens up super-speed, low-latency connectivity to an exponentially growing fleet of connected cars, machines, appliances, and gadgets.

## The Edge Computing Renaissance with a Halo of Cloud Computing

Central offices<sup>2</sup>, along with cable MSO headends, are brains of a service provider's nervous system and they are evolving into something resembling small cloud data centers as virtual servers, software-defined networks (SDN), and network function virtualization (NFV) have replaced fixed-function hardware devices. As SPs prepare for an onslaught of high-speed devices and bandwidth-hogging content, they face the limits of inflexible, hardware-based centralization and the need to redistribute workloads to be near the demand. The force behind SP infrastructure disaggregation is an overarching need for scalable performance driven by the following factors:

- Although telecom providers have already deployed **5G service** in select markets, the mass adoption of 5G devices will take years. From less than 1 percent of devices shipped this year, **IDC estimates** that 5G phones will make up more than a quarter of the market by 2023. However, SPs won't be able to deliver on 5G's promise of millisecond latency with 1 Gbps or higher download speeds without offloading backbone networks and moving content and services closer to the user.

---

<sup>1</sup> The quote is attributed to him though provenance is unverified.

<sup>2</sup> We'll use the general term Next-Gen Central Office to refer to actual central office locations as well as cable headends for multiservice operators and any edge-located mini data center that functions as a multi-tenant, application-hosting compute, storage, and networking platform.

- Changing consumer behavior inside the home has stressed existing SP infrastructure. The **disaggregation of the cable bundle**, via increased use of OTT streaming services and concomitant reduction in broadcast or shared cable media, has significantly increased the demand on last-mile network capacity.
- Beyond residential services, the rising strategic importance of the CO and MSO headend locations, in providing **support for enterprise SD-WAN** and edge services, requires that the next generation of these locations provide a robust foundation for hosting new services that could supplement or replace existing customer premise equipment (CPE) capabilities. These locations could also act as the on-ramps for express lane access to public cloud infrastructure providing infrastructure and application services.
- The **virtualization of carrier networks** from fixed function hardware to SDN and NFV has improved the efficiency and capacity utilization of carrier infrastructure and their ability to rapidly introduce new services. These improvements have also increased customer expectations and the need for service innovation and increased speed and flexibility in service delivery. However, software-defined infrastructure also facilitates moving services to next-generation platforms like containers, serverless functions, and edge data centers to create competitive differentiation.
- While the Internet might have **flattened the world, a countervailing increase in the diversity of regulatory requirements across different countries** has partially Balkanized the global network and made compliance problematic, especially when services are delivered from central locations.
- The **growth in IoT usage** of all types – intelligent sensors, connected cars and smart city infrastructure among others – increases the demands on capacity and further diversifies traffic sources and last mile traffic characteristics. Furthermore, future IoT endpoints will rely on the low latency promised by 5G service, something that is impossible to deliver without localized service delivery. Indeed, the need for fast, low-latency 5G connections is imperative for future autonomous vehicles and ADAS systems that might use some form of peer-to-peer or hyper-local network for intervehicle communication.


**As SPs prepare for an onslaught of high-speed devices and bandwidth-hogging content, they face the limits of inflexible, hardware-based centralization and the need to redistribute workloads to be near the demand.**

## The Changing Carrier Infrastructure Landscape

5G is the most significant factor affecting SP infrastructure, however it's like a hurricane in that providers have ample warning of the changing environment. **Deloitte estimates** that 72 operators were testing 5G in 2018; 25 will launch partial service this year and another 26 in 2020. Fortunately, initial fears of 5G blowing the top off of carrier CapEx budgets appear to be unfounded. As Deloitte notes, after performing field trials, operators told investors to expect roughly flat CapEx spending relative to the 4G rollout. The factors behind such a significant change are unclear, however one might be consequent plans to obviate the need for massive fiber and central office expansion by delivering more services from minicloud infrastructure in base stations and other edge locations.

### Only one in seven mobile connections will be 5G by 2025

Global mobile adoption by technology, share of mobile connections, excluding cellular IoT



Source: GSMA, *The mobile economy*, 2018.

For 5G edge deployments, often called **multi-access edge computing or MEC**, the next generation resembles a scaled down CO that is positioned as close as feasible to users, ideally within one hop. These MEC systems often rely on a carrier’s core cloud infrastructure for services and content that can’t efficiently be locally delivered.

While estimates for next-generation central office and carrier edge spending are hard to come by, **various studies show** a flourishing market for edge infrastructure writ large with spending growing at double-digit annual rates. Bolstering the case for distributed carrier infrastructure, Gartner says that latency and/or bandwidth will be critical factors for 30 percent of all application workloads by 2021.

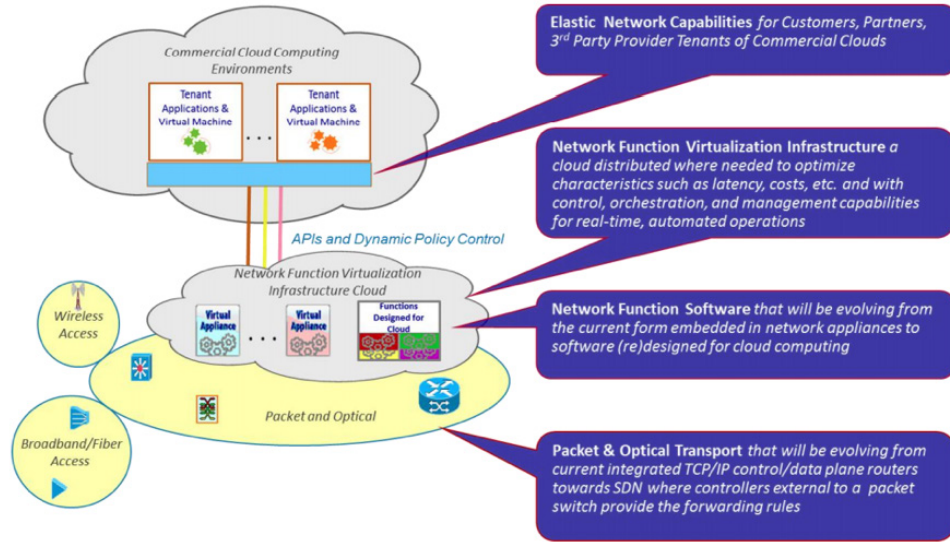
### NGCO Characteristics and Technologies

Our discussion of the causes and enablers of next-generation central offices and future carrier infrastructure provides a hint at the capabilities, characteristics, and technologies that make up NGCOs, but we’ll enumerate them in more detail here. Although the term “next generation” makes such network designs seem futuristic, the concepts are a decade old. Indeed, AT&T encapsulated many of the fundamental tenets of an NGCO more than six years ago when it laid out its **Domain 2.0 vision**:

*“In a nutshell, Domain 2.0 seeks to transform AT&T’s networking businesses from their current state to a future state where they are **provided in a manner very similar to cloud computing services**. ...*

*The replacement technology consists of a **substrate of networking capability, often called Network Function Virtualization Infrastructure (NFVI) or simply infrastructure that is capable of being directed with software and Software Defined Networking (SDN) protocols** to perform a broad variety of network functions and services. This infrastructure is expected to be comprised of several types of substrate. The most typical type of **substrate being servers that support NFV, followed by packet forwarding capabilities based on merchant silicon, which we often call white boxes**.*

***AT&T services will increasingly become cloud-centric workloads.** Starting in data centers (DC) and **at the network edges** – networking services, capabilities, and business policies will be instantiated as needed over the aforementioned common infrastructure.”*



**Figure 1 – High Level Cloud Networking Architecture**

Source: AT&T Domain 2.0 Vision White Paper

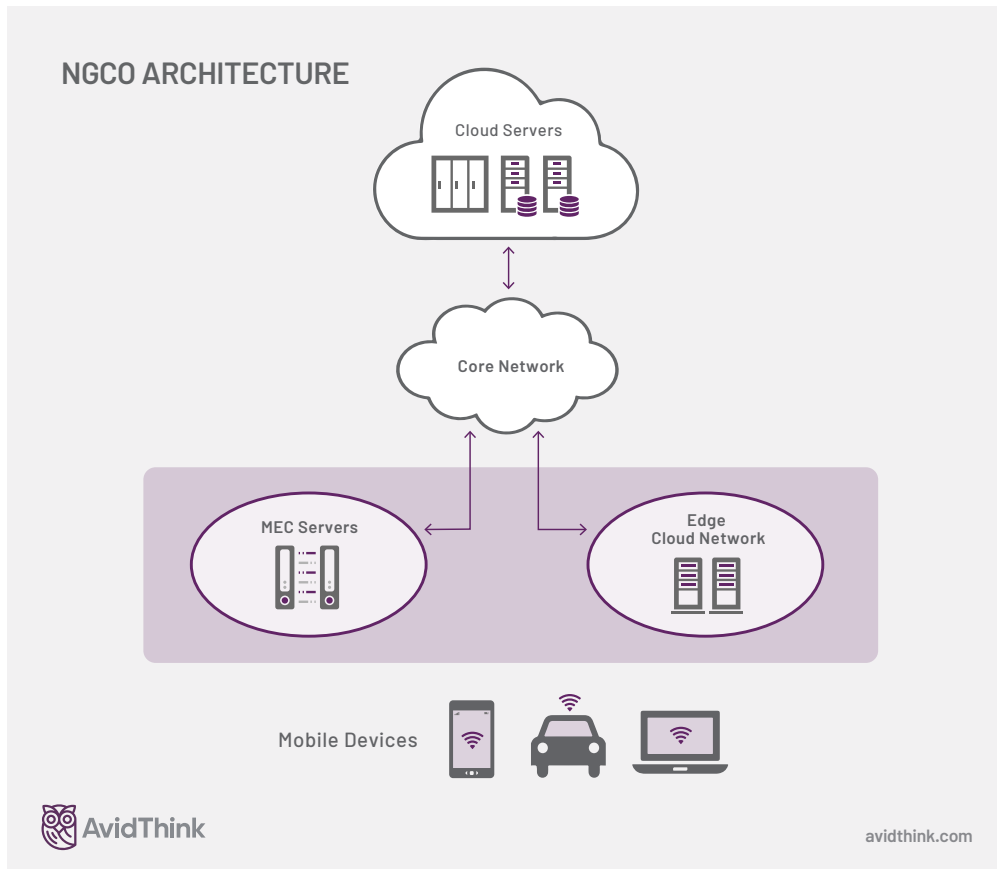
The fundamental characteristics of the NGCO are enabled by several technologies that can transform it into a fully virtualized, cloud-like environment for network service operations and delivery. These technologies are:

- **SDN** with separate data and control planes delivering the software abstractions that allow decoupling network configuration and packet control from the hardware implementation. Furthermore, turning network transport into a software layer facilitates greater automation of administrative tasks and the programmatic control of network parameters based on real-time measurements and events.
- **NFVI and virtual network functions (VNF)** that replace the fixed-function hardware appliances with adaptable, scalable software services. NFVI itself resembles cloud infrastructure by using commodity (usually x86) servers with virtualization stacks like **OpenStack** or **VMware vCloud** as the computing substrate. VNFs are the L4-7 applications that may take advantage of optimizations via APIs exposed by the infrastructure and service layer substrate.
- **High-density, rack-mounted commodity white box or OEM servers and switches** that make up the NFVI. This is not different from most cloud infrastructure. However, when deployed at edge locations, the systems must be suitable for thermally- and (sometimes) power-constrained environments and occasionally, harsh conditions.
- A **virtual infrastructure management system** along with a **service orchestrator (NFVO)** that connects to the NFVI manager and controls instances, workload scheduling, monitoring, and the service catalog.

Despite its lackluster success in the enterprise market, OpenStack is still a popular platform for NFVI due to being an open source project and its ability to be customized and easily integrated with other software components in an NGCO. Many SPs see the virtualization of the central office hardware and software stack as an opportunity to embrace open source, an inclination that has spawned several related open source projects.

Increasingly, NFVI are supporting Linux containers, as lower-overhead, more agile platform for VNFs—either directly on bare metal, or within lightweight VMs. In addition, these container-based VNFs are delivered using a cluster of nodes and container orchestration software like Kubernetes to improve resource utilization and service scalability. Implementing VNFs as containers also facilitates lightweight service chaining in which several VNFs are connected to create a composite service.

Borrowing from microservice application designs, service chains promote reusing and linking simpler VNFs that deliver a particular feature to create more advanced services. As part of this architectural evolution, many of these VNFs are built using microservices architecture that employs a **service mesh** such as Istio/Envoy for communication and coordination. Whether the move to containers results in a complete move away from OpenStack – or vCloud – or whether these containers are integrated into an OpenStack managed infrastructure remains to be seen.



In sum, **an NGCO architecture** uses scale out, commodity (sometimes ruggedized) hardware to build a virtual, cloud infrastructure, in which VMs and containers run software-abstracted services that are managed and orchestrated by a central system and exposed and consumed through APIs.

## Hardware Components and Design

As we've seen, an NGCO looks much like conventional cloud infrastructure, both logically and physically, though in some locations they likely will be ruggedized servers. As with the large cloud builders, SPs have embraced high-density commodity servers and edge (ToR) switches that can be bought in bulk, deployed across many different edge locations, and used to power a variety of VNFs.

There are two noteworthy physical configurations for servers:

1. **HCI systems** from OEMs like Supermicro, Inspur, Quanta, and others that are functionally identical to enterprise HCI products from Dell EMC, HP Enterprise, and Nutanix that bundle a virtualization system. The hardware comes in a

variety of form factors, for example 1U servers might be all that's needed for some edge locations, while larger edge sites might need the capacity and density of a 2U system supporting four half-width server blades.

**2. Rack-scale systems** using Open Compute Project (OCP) standards that are available from many of the OEMs building HCI products. Although OCP is primarily focused on hyperscale cloud installations, **some vendors like Nokia use the spec for OCP equipment tailored to telco** central offices. Smaller OCP configurations, for example using redundant server chassis, storage trays, switches and power shelves occupying a half rack might be appropriate for larger edge locations such as 5G base stations or a central hub for an office park.

NFVI has become such a significant market that system builders and component suppliers now sell products designed for networking applications. For example, Intel's **second-generation Xeon Scalable processors (Cascade Lake) are available in N SKUs** that deliver a higher base clock frequency and the ability to independently increase the clock speed on a few cores to accelerate single-threaded communications workloads. Intel benchmarks these models as 58-76 percent faster than first-generation (Skylake) scalable models for typical VNF workloads, such as **OVS packet processing using the DPDK and vector packet processing (VPP)**.

Smaller edge locations in thermally and power-constrained environments need components and systems that are designed to optimize performance within the physical limits of an edge location. For example, NFVI servers might opt for **Intel's Xeon D processors** designed for edge network, control plane, and mid-range storage systems that combine up to 8 cores with integrated memory and PCIe controllers, Ethernet, and other I/O interfaces on a **single SoC**. Importantly for edge deployments, the D Series parts can provide four Broadwell (5th-generation) cores with four 10 GbE interfaces in a 35 watt package. The **recent D-1600 series** delivers 8-25 percent more performance for networking workloads than the prior generation within the same power and cost envelope.

SDN and NFV offer NGCOs great flexibility in adapting to changing workloads and requirements, but operators can still run into limitations to the degree that some network functions are implemented by hardware Application Specific Integrated Circuits (ASICs). Often, the processing of data plane packets is hardcoded, which hampers adding network features that might be required for new IoT or security protocols, since it means refreshing the hardware. This rigidity can be mitigated by making the network devices more programmable, such as **making the data plane processing itself programmable by using FPGAs** that can be reprogrammed after deployment.

**Performance limitations of general-purpose CPUs in these edge locations can be addressed through the use of SmartNICs based usually on FPGA, Network Processor Units (NPU) or ASICs.**

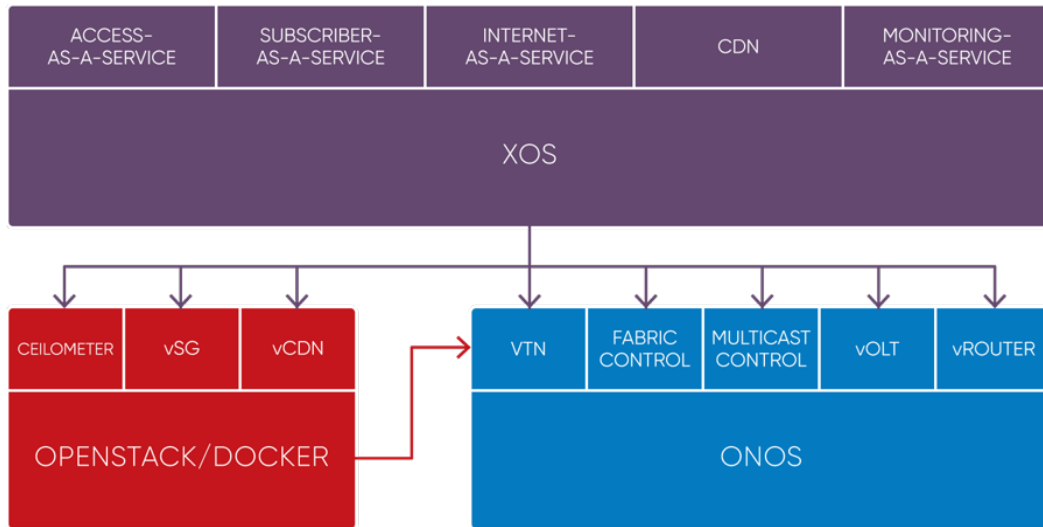
Further, performance limitations of general-purpose CPUs in these edge locations can be addressed through the use of SmartNICs based usually on FPGA, Network Processor Units (NPU) or ASICs. Vendors such as Ethernity, Mellanox (part of nVidia), Napatech, Netcope, Netronome, and SolarFlare are pushing to have SmartNICs included as part of standard server platforms to offload the CPU for I/O-heavy workloads. These add-on cards can enable white box servers to rival the performance of traditional hardware appliances and should see increased use by SPs in both central and edge locations. From a price-performance basis, the premise is that CPU cycles are precious resources, especially at the

edge, and should be dedicated to processing workloads that require their flexibility, while the SmartNICs are assigned to perform the more deterministic workloads of processing and handling network packets.

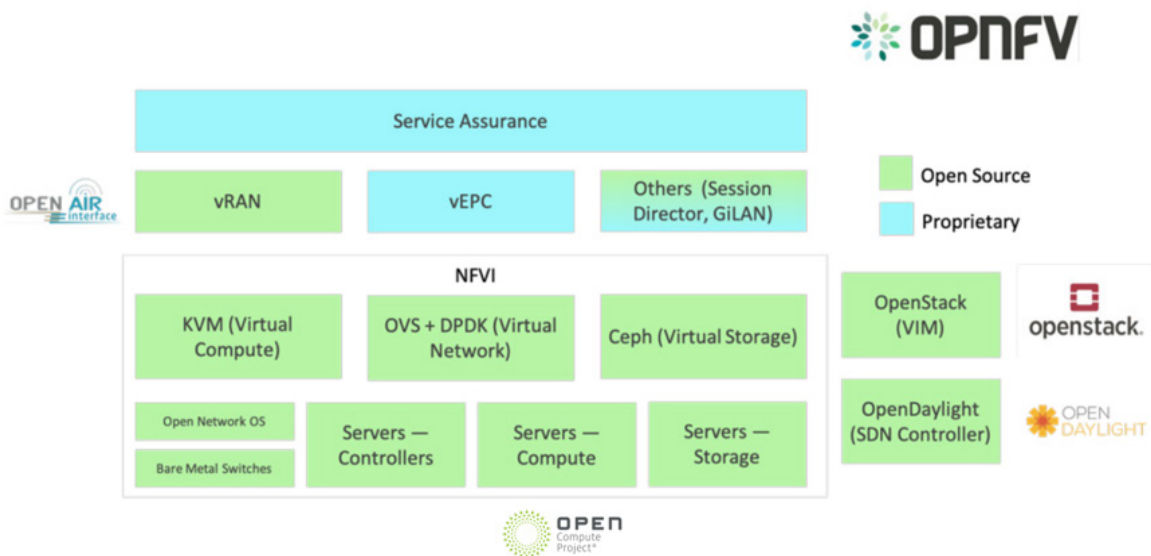
## NGCO Software Foundation and Attributes

Looking more closely at the software side of NGCOs, we see a multilayered stack built on a foundation of cloud virtualization and management platforms and SDN and NFV software. To this SP-specific features are added through an extensive service platform.

There are a few different software stacks being developed for the NGCO. Two of the most prominent are the Open Networking Foundation's (ONF) Central Office Rearchitected as a Datacenter (CORD) and the Linux Foundation's Open Platform for NFV's (OPNFV) Virtual Central Office (VCO). CORD has been incorporated into solutions for the mobile and residential markets in the form of the Converged Multi-Access and Core (COMAC) and SDN-Enabled Broadband Access (SEBA) projects under the ONF.



Source: [ONF CORD Platform](#)



Source: [OPNFV VCO White Paper](#)

For both stacks, and likely true for any other NGCO software stack, there are the following key components:

- **Cloud virtual infrastructure stack.** Examples are OpenStack with KVM, VMware with ESXi or bare metal infrastructure running Docker containers.
- **SDN software stack.** Usually an SDN Controller such as ONF's Open Source Network Operating System (ONOS) or the Linux Foundation's OpenDaylight, or Open Kilda and others. They provide the software abstractions and APIs enabling north- and south-bound network interfaces for features such as packet switching/forwarding, flow monitoring, network configuration and plug-in applications.
- **Service control plane and VNF application platform.** In CORD, this role is played by **ONF XOS** or a commercial alternative like **VMware vCloud NFV could be used**. This usually plays the role of a VNF manager and provides orchestration as well.
- **Monitoring and assurance.** There will generally be a layer that provides service assurance and monitoring to ensure the overall system health and may employ self-healing measures to ensure uptime, such as restarting unresponsive VNFs or VNF components.

In a situation where containers are used, a **Container orchestration stack**, typically Kubernetes, will also be employed to deploy, manage, and scale network service containers.

There will also be a need for higher-level orchestration across multiple services, for which a project like **Open Network Automation Platform (ONAP)** from the Linux Foundation could be a fit. Further, in production, the overall stack needs to be integrated with a service provider's **operations and billing support systems (OSS/BSS)**, which are used for customer ordering, service provisioning, billing, network inventory, and problem management. These are typically proprietary platforms unique to each SP and therefore a custom integration effort.

Beyond CORD and VCO, SPs might be interesting in exploring **OpenNESS**, a software framework seeded by Intel, to develop and deploy network applications at the edge or on premises by creating an abstraction layer that exposes standards-based APIs from 3GPP and ETSI MEC industry groups to application developers.

**A software-defined infrastructure that can span both massive scale out COs and small edge locations offers carriers and other network service providers significant financial, technical, operational and competitive benefits.**

## Benefits of NGCOs

The NGCO is a major architectural departure from traditional telco designs that requires both new hardware and software. As such, building them can't be a one-off project in just a few locations, nor can it be undertaken lightly. Nonetheless, a software-defined infrastructure that can span both massive scale out COs and small edge locations offers carriers and other network service providers significant financial, technical, operational and competitive benefits, most notably:

- The **architectural decoupling** of data and control planes, which leads to a concomitant decoupling of service delivery from particular pieces of hardware and that allows for **independently scaling capacity and services/content**.
- **More flexibility in service customization**, including the ability to **deliver customized services and applications from a microcloud** targeted to particular customers and with contextualized services and resource allocations tailored to differing needs within these groups.
- **More choice of hardware and software products** and suppliers due to using industry-standard commodity hardware and OSS, which conversely leads to less **equipment diversity**, and **fewer technical support challenges**, since many components are interchangeable and the scalable design enables greater resource consolidation where it makes sense.

- **Faster time to market** for new services and updates, including all-important security patches, since they are implemented through software. Further the use of standard APIs and containerized microservice architectures can facilitate re-use of components.
- **Greater scalability** by replacing physical hardware with virtual services and that allows for the rapid expansion of service capacity in response to short-term spikes in demand.
- A **cheaper implementation over the long term** by reducing both CapEx and OpEx. It is important to note that long-term gain doesn't come without short-term pains, since NGCOs implementations will face startup costs for new software and equipment, system integration, customization, and testing of commodity and OSS components. Costs could be held in check by using commodity hardware and open source software and via increased automation of the virtual control plane.
- **Reduced traffic and congestion in the network core** due to edge COs offloading high-demand content and services, which leads to lower latency for service and content delivery of high-demand services.
- Theoretical **improvements in system-wide energy efficiency** due to fewer required network hops in scenarios using edge COs and hence fewer pieces of equipment being used to deliver a particular transaction.

On top of these competitive and operational benefits, a software-defined NGCO has strategic value by providing the foundation for secure, multitenant services that enable new business models such as the ability to more easily resell capacity and services to third-party service providers; provide service bundles that include third-party products; and allow customer self-service installation and management of select network services.

## Recommendations and Next Steps for Service Providers and Enterprises

The design, implementation and testing of software-defined NGCOs based on open hardware and software is a multiyear endeavor that can only be undertaken with sponsorship at the highest levels of an organization. The first step involves taking the vision we've outlined in this paper and applying it to the unique circumstances of an organization's infrastructure, technology expertise, and customer needs.

Next, for SPs who have not yet embarked on the SDN and NFV journey, they should get started by applying server and network virtualization in opportune situations such as SD-WAN, including the use of multifunction servers instead of hardware appliances for CPE. Over time, an organization must plan how to gradually virtualize the entire network, both the control plane and network services since a next-generation edge CO makes little sense when it's connected to a core made up of legacy technology.

Other steps on the path to NGCOs include:

- Getting involved with open source projects we've covered above that are building the technology foundation for SDN, NFV and NGCOs.
- Developing usage scenarios that can exploit edge service delivery and where moving services and data to the edge provide the most value in better performance, reduced congestion, and higher customer satisfaction.

Building an NGCO that includes distributed edge sites is a strategic decision that entails a lengthy project and significant upfront costs. While the design offers long-term financial and competitive benefits, the startup costs mean that NGCOs should only be undertaken when an organization is already facing a major system and equipment upgrade cycle. Fortunately, the ongoing 5G wireless transition, the increased use of consumer OTT streaming services, and an overall explosion of enterprise data makes now the perfect moment to take the plunge into a software-defined future.



**AvidThink, LLC**  
1900 Camden Ave  
San Jose, California 95124 USA  
avidthink.com

© Copyright 2019 AvidThinkI, LLC, All Rights Reserved  
This material may not be copied, reproduced, or modified in whole or in part for any purpose except with express written permission from an authorized representative of AvidThink, LLC. In addition to such written permission to copy, reproduce, or modify this document in whole or part, an acknowledgement of the authors of the document and all applicable portions of the copyright notice must be clearly referenced. All Rights Reserved.