



Who dares wins!

How access transformation can fast-track evolution of operator production platforms

September, 2019

Arthur D Little



T . . *Telefonica*

Content

Summary and key ideas	3
How to read this document	4
Introduction	5
1. Why business as usual must change	7
2. The central office pod – <i>A winning cloud design</i>	10
3. The case for the central office pod	14
4. Leading the creation of a <i>carrier-cloud</i> mind-set	19
5. Launching a <i>minimum viable</i> access transformation program	22
6. What's new about the <i>new design</i>	27
7. How the central office pod changes everything	33
8. Operationalizing the central office pod	36
9. Building service and platform engineering capabilities	39
10. The technologist's guide to getting started	42
Closing remarks	47
Glossary	51

Authors:

Arthur D. Little: Salman Ali, Mariana Atilano, Carlos Mira, Jesús Portal

AT&T: Tom Anschutz, Mitch Olson, Earl Pope

Deutsche Telekom: Bodo Jacobs, Hans-Joerg Kolbe, Manuel Paul, Fabian Schneider, Robert Soukup

Telefónica: David Artuñedo, Alfonso Carrillo, Patrick López, David López Meco

Safe harbor statement

Arthur D. Little, AT&T, DT and Telefonica ("the parties") are committed to adhering strictly to the letter and spirit of the antitrust and securities laws. Any and all meetings, communications, conference calls, shared resources and collaborative workspaces (including, without limitation, documents, emails, presentations and draft notes) conducted under the auspices of this collaboration (collectively referred to as the "activities") are designed solely to provide a forum or means for the expression of various points of view. Under no circumstances should the activities be interpreted as activities used as a means for competing companies to reach any understanding, expressed or implied, to restrict competition, or in any way impair the ability of members to exercise independent business judgment regarding matters affecting competition. All information, data and analysis set forth in this report, including economic estimates and statements as to the effects, expected timing, readiness of certain technologies and methods, may constitute forward-looking statements within the meaning of the safe harbor provisions of the Private Securities Litigation Reform Act of 1995. Such estimates and statements include, but are not limited to, statements about the objectives, benefits of the technology, including future revenues and efficiency benefits as well as expectations and intentions, and other statements that are not historical facts. Such economic analyses and statements have been developed by Arthur D. Little in their entirety based on their experience and beliefs and expectations and are subject to significant risks and uncertainties. While the parties have reviewed the work, we expressly disclaim that these estimates or statements represent our specific approach, strategy or plans. Neither should this be understood to mean that the parties are under any obligation to report how they are adopting these technologies and methods, or to update, alter, or otherwise revise any past forward-looking statements regarding our strategy, whether written or oral. Persons reading this report are cautioned not to place any reliance on these forward-looking statements, which speak only as of the date hereof. Among the risks and uncertainties that could cause actual results to differ from those described in the forward-looking statements are the following: (1) the occurrence of any event, change or other circumstances that could give rise to previously unidentified technical risks, (2) risks that technical prerequisites may not be satisfied in a timely manner, (3) risks that sufficient operators will not support new direction, (4) the risk that our suppliers or others may not be willing to adopt the proposed technical architectures, (4) the risk that regulatory approvals may not be obtained or may be obtained subject to conditions that are not anticipated, (6) the ability to attract, retain new personnel and (7) failure to realize the benefits from adoption of the technology.

Rights and permissions

This work is made available under the Creative Commons Attribution 4.0 International license (CC BY 4.0) <https://creativecommons.org/licenses/by/4.0>



Summary and key ideas

Rapid advances in data center and cloud-based technologies have put the entire operator ecosystem in motion. Operators and their shareholders want to use the automation advantages of deploying and managing infrastructure and services with modern cloud tooling, but have a large, embedded base of legacy OSSs and BSSs to deal with. Not all operators will change, or can adapt, at the same rate as this technological and cultural disruption, so industry norms and standards bodies are struggling to provide timely and usable solutions to the hard problems facing the industry. Traditional telecom equipment vendors are in a predicament. Their customers want solutions with cloud-like simplicity, open interoperability and performance, but shifting to open architectures puts their business models at risk. They'll need a credible alternative that will allow the industry to modernize, as well as regenerate the supply chain.

Recognizing the challenge, a handful of operators, including AT&T, Deutsche Telekom and Telefónica, are defining an alternate pathway. Their pioneering approach relies on disaggregating and virtualizing the access network to considerably expand the role of the central office, in order to include fixed and mobile aggregation and edge cloud services. This paper describes their ideas and provides insight into their real-life experiences to make the new approach production ready. While numerous execution challenges require further attention, this paper makes the case that, for "those who dare," access-driven transformation provides an exciting alternative to the status quo.

Key messages of this report

1. Operators shouldn't wait to help themselves transform.
2. To remain relevant, the industry must adopt data center technologies and align operating models with cloud-based ways of working.
3. The "CO pod" provides operators with a new toolbox to virtualize/re-engineer existing services, as well as prototype and test new service ideas, based on cloud-hardened development and operations methodologies.
4. Operators must decide if they will lead or follow; for "those who dare," the prize could be significant.
5. Strategic priorities, not technology, must drive access-driven transformation programs.
6. Access-driven transformation replaces the traditional central office aggregation function with a leaner and lower-cost design that reduces TTM.
7. The CO pod considerably widens choices for operator production platforms.
8. Making the new design a reality requires operators to address four operationalization challenges.
9. Success requires operators to acquire new skills in engineering and supply chain, and to rethink investment planning.
10. There is no "best approach": program, design and monitoring must be tightly linked to corporate KPIs.

How to read this document

This document is divided into two modules. The first module is directed at a generalist and technical audience, targeting operator decision-makers from corporate development, as well as product managers, marketers and C-levels. The second module is a deep-dive directed at an operational and technical audience and is intended to encourage greater investment and resource dedication to access-driven transformation, as well as foster community among operators.

The central office pod for generalists

This generalist module is contained in the five sections. The first section outlines why existing operator transformation plans are challenged. The second section describes how access-enabled transformation can enable an alternate production platform and product development pathways. The third section examines the strategic option space enabled by the CO pod, quantifying the key value drivers for change. The fourth section looks at the leadership challenges associated with pivoting toward a cloud-based business and makes the case to upgrade operator skills and expand the supplier ecosystem. Finally, the fifth sections contains insight from AT&T, Deutsche Telekom and Telefónica on how to structure and launch a minimum viable program.

The central office pod for technologists

The technical deep dive module (background in blue) is directed at a technical audience focused on network architecture, operations and supply chain management, as well as existing and prospective operator technology suppliers. It is contained in five additional sections beyond the five generalist sections outlined above. Starting from section six, we outline the design for new edge aggregation and a cloud pod that underpins the CO pod. Section seven examines how the new design creates new options to change the operator production platform. The eighth section examines the operational implications of onboarding the new design. The ninth section looks at how to rebuild core capabilities and supplier ecosystems. Finally, the tenth section contains recommendations for how to get started, drawing on lessons learned from AT&T, Deutsche Telekom and Telefónica.

Acknowledgments

This paper has benefited from the collaborative effort of many individuals who have been involved: from AT&T, Deutsche Telekom and Telefonica, as well as Arthur D. Little. Without their support, dedication and passion, this work would not have been as enriching or, indeed, possible. This paper is a tribute to their generosity and patience, and their passion to contribute to creating a thrilling future.

Introduction

Three powerful forces are changing the telecoms industry. The first is the well-understood concept of *demand growing faster than revenues*. Second is the macro-trend of *technological convergence* pushing “open” data center technologies, tools and techniques as alternatives to appliances, based on proprietary technology, to improve business velocity and automation. Third is the increasing amount of value being generated by third parties, as well as the operator desire to integrate this value into existing businesses and participate in these markets. These forces are pushing the industry to align or converge with others: the industry must accelerate the transformation of its production model, adopting, where practical, cloud technologies and aligning its operating model with cloud ways of working. However, this is easier said than done. There is not yet a proven approach or credible supply chain to bring the cloud into the operator network, and one may not emerge until the industry takes back control of its technical destiny.

This white paper describes an ambitious or *moonshot-like* approach to operator transformation. Such an approach is needed because the authors feel the current industry approach or *business as usual* cannot deliver on the industry needs, as they are demanded today. The focal point of the new concept is deliberate alignment with cloud technologies and operations, which lead to re-architected access networks that can support more than just access services. It is based on lessons learned from traditional modernization programs, as well as emerging practices in disaggregation, virtualization, employing common data center equipment, softwarization and simplification of the telco stack, and closed-loop IT automation from the cloud. The new approach combines business and technical thought leadership from AT&T, Deutsche Telekom and Telefónica, as well as strategic and economic insight from Arthur D. Little. Collectively we are promulgating a new design for the operator production platform: one that allows the industry to take back control of technology, attract urgently needed talent, re-skill the workforce and diversify the supplier ecosystem.

The innovation behind the new design is a modular Central Office pod¹ (“CO pod”) architecture, designed to replace traditional access network technologies. It is inspired by Central Office Re-architected as a Datacenter (“CORD”) from ON.Lab (now ONF²), as well as AT&T’s, Deutsche Telekom’s, and Telefónica’s experiences with Domain 2.0, DTBen and Unic@ programs, respectively. The CO pod disaggregates proprietary vertical equipment, like Multi-Service Access Nodes (MSANs), into general-purpose hardware and reproduces the relevant functions in software managed by standard IT automation tools. This approach introduces well-understood efficiency gains through infrastructure pooling and reuse. However, it also puts the industry on a more sustainable pathway. Disaggregating proprietary equipment opens up access hardware and software to new suppliers, integrators and convergence opportunities with adjacent telecom and non-telecom uses. Use of cloud practices and open source accelerates innovation and lowers total cost of ownership. Adoption of de facto standard IT tooling and operations increases the addressable skilled-labor pool. Moreover, disaggregation affords the opportunity to rethink existing BSS/omnichannel platform architectures and addresses new business value pools through clever use of edge infrastructure.

The CO pod is not a concept or product that belongs to someone. It is a “community-developed” (see inset: “[What is community?](#)”) template for enabling broadband and mobile access to use SDN and cloud operational paradigms, where the coauthors of this paper are major contributors and thought leaders. The community emerged as a shared interest group to help close the gap between telecoms and cloud architectural thinking. The community did and continues to learn much from adjacent “IT” and web-scale industries. However, while many of the concepts and approaches are readily applicable to telecom problems, the community recognized the designs cannot be adopted verbatim. A proven set of templates that directly address industry challenges is needed. These include platform engineering to ensure carrier-grade performance, building the right tools to enable services innovation, and finding the right balance between technology control and vendor dependence.

¹ In this context, CO pod refers to a module of network, compute, storage, and application components that work together to deliver networking services. It is a repeatable design pattern, and its components maximize the scalability, and manageability

² Open Networking Foundation

The authors are not alone on this journey; under the auspices of ONF, there are multiple³ multidisciplinary teams, working on redefining the industry using the CO pod architecture. This linkage between IT and telecoms is not a technical fork, but rather, technological convergence or new co-reliance between telecoms and these adjacent industries and it is important for decision makers to understand that continuous engagement with others that espouse and invest in maturing virtualization and other cloud technologies is vital.

It is the collective view of the authors that the CO pod provides a new toolset for broader operator transformation, as well as a platform to respond to the three forces described at the start of this paper. Enabled by its front-line location, it provides operators with a safe place⁴ to prototype and test new service ideas, as well as support existing vetted services using well-understood cloud development and operations methodologies. Moreover, distributed deployment means more resilient services and a smaller span of outages by following well understood principles of cloud engineering.

However, there is no such thing as a free lunch. Deploying at the edge requires touching many more locations than when deploying in fewer, more centralized locations. Gaining the benefits of cloud technologies and practices requires strong leadership as well as unwavering commitment to operational and organizational change. Operators must learn the art of using general-purpose hardware, network programmability, and virtualization software to produce carrier-grade products and services in which software architecture provides high availability. In this future, the new norm will be “do it together”: combining one or more pieces of open-source and commercial software with functionality provided by the CO pod to create new applications. This means operators must become more self-reliant and take greater responsibility in the value chain to define, develop, integrate, test and commission in-house solutions to drive differentiation. This change will require leadership not only from technology, but also from business and commercial functions, who must take advantage of the technology to rethink how to compete. These ideas present numerous execution challenges, but, for “*those who dare*”, it provides a viable alternative to existing production approaches, as well as the tools to carve out a role in the business of cloud.

What is community?

Each operator does not have to go its own way; it is possible to collaborate, or form communities emulating collaborated development of open-source software.

Recognizing the change is a common challenge facing the industry; operators can pool technical and financial resources and direct them toward creating one or more communities. In these communities, operators can define the rules of engagement to ensure the industry gets what it needs. It also creates opportunities for others to build new businesses and business models. This is the essence of a community-based approach: a number of operators working collaboratively with a shared ambition, focusing on common elements that do not drive differentiation. The more communities, the larger they are, and the wider the community breadth of expertise is, the easier, better, and faster it is for all involved.

This is not something new; numerous operators are already active in communities closely related to telecommunications, such as ONAP, OCP, OEC Forum, TIP, and ONF.

To socialize the new architectural concepts and implications, the paper is structured in two modules: Module 1 is directed at both generalist and technical audiences; Whereas Module 2 is a technical deep-dive aimed at a predominantly technical audience.

Together they are intended to encourage greater collaboration and investment in technical and organizational development based on the *new design*. In sharing this insight, AT&T, Deutsche Telekom and Telefónica intend to socialize the CO pod concepts, generate wider interest, and ultimately develop a larger community to lead the process of proving, hardening, and turning it into a *de facto* standard. For the CO pod to become a reality more quickly, it will require a concerted effort by many in the industry, not just a few.

³ Key ONF operator members include China Unicom, Comcast, Google, NTT, Turk Telecom as well as AT&T, Deutsche Telekom and Telefónica

⁴ This is in sharp contrast to traditional services that need to be deployed in the core, where failure could affect millions of users

1. Why business as usual must change

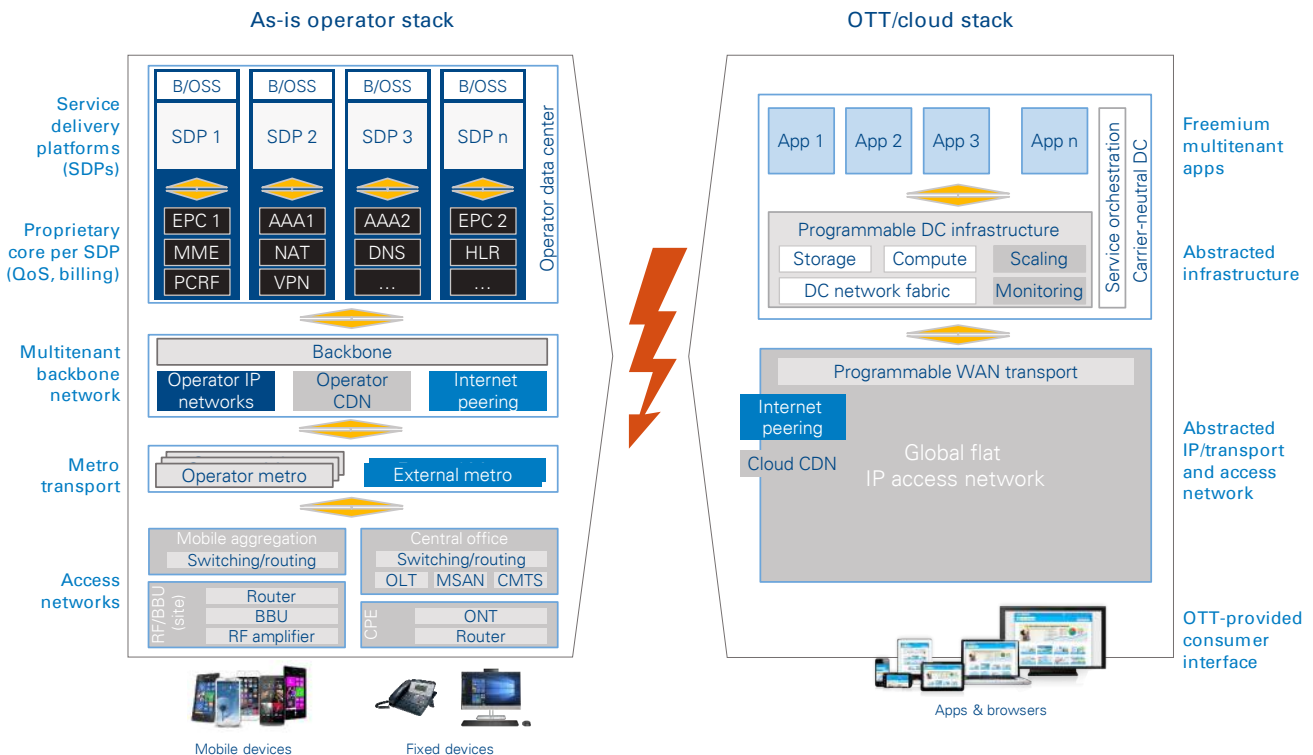
Operators shouldn't wait to help themselves transform:

- The cloud architectural pattern is the focus of innovation in multiple industries – and drives technological convergence across all of them.
- Operators need to strengthen their technical and operational capabilities to onboard new technologies and innovate; this is evidenced by initial SDN and NFV deployments' inability to deliver real benefits to operators.
- Operators should not expect standards development organizations (SDOs) or existing vendors to deliver a credible pathway to the cloud.
- Operators must espouse cloud technology, operations and ecosystems, in order to help themselves.

A substantial cloud ecosystem has emerged, which serves both consumer and business users and relies on access connectivity provided by operators around the globe. It consists of a wide range of companies that are described with terms such as "over the top" (OTT), cloud service providers, and web-scale operators. They provide a diverse range of services, from infrastructure to platform and application services. Despite their diversity, they all have one thing in common: they are unlike traditional telecoms operators. As shown in Figure 1, the scope

of their activities and operating models is based on a very different technology architecture. OTTs typically provide services built on virtualized, abstracted and programmable compute, storage and networking resource pools whose underlying technology is a multitenant platform, accessible on demand to everyone connected to the World Wide Web. In this ecosystem, the only barriers to entry for a new global OTT are software development skills and a credit card. Success has propelled some cloud companies much further, building hyper-scale

Figure 1: Comparison of the cloud and the operator stack



infrastructure and ecosystems around search engines, social media, video distribution, infrastructure and devices. Despite the potential regulatory and security shortcomings of this new cloud ecosystem, it has become the *force du jour*, upending the dominance of the traditional IT and private networking stack. It has attracted services that once were siloed to operate over the global Internet, such as VoIP, IPTV and instant messaging. It has also created many more (myriad) services and capabilities that it would have been inconceivable to create any other way. For operators, sandwiched between OTTs and their customers, the innovation and traffic-growth cycle means operators must continuously upgrade access networks to cope with increasing traffic from customers loyal to their smartphones, tablets and apps. This is because, despite their excitement about cloud services, consumers' and businesses' willingness to pay operators for incremental traffic is null or very low.

Few operators would claim mastery of the cloud architectural pattern applied to communications. Beyond internal culture and established architectures and practices, operators are reliant on standards development organizations (SDOs) and a shrinking set of global vendors, which have been slow to provide solutions for operators in order to take advantage of the technologies and business practices in the adjacent cloud markets.

The origin of the problem is how the industry innovates. Future industry offerings are based on the work of SDOs that collectively determine the industry roadmap, setting the overall direction and pace (see inset: "System standards"). Their well-intentioned interoperability- and compatibility-focused work is developed through a gradual process of consensus building among operators and their suppliers. The outcome is often complex and detailed guidelines and specifications that attempt to balance the conflicting interests of the parties, rather than emphasize innovation.

The existing innovation model has allowed the industry to achieve many things, such as xDSL and FTTx, and to go from GSM to 5G. However, there are several unfortunate aspects and effects of this standards-driven process. It is inward-looking and, consequently, often late to adopt new thinking. It's especially difficult to provide reasonable standards when technology jumps from one adoption curve to the next. To date, few SDOs have been able to internalize cloud technologies and methods or shape the operator service portfolio toward the cloud era.

System standards

The role of standards is not all good or all bad. There is clearly a place for, and benefits from, many types of standards. Much business and technology innovation stems from basic or "component" standards that support global interoperability and reuse of common components across sets of systems, including operator networks. Examples of component standards include ethernet, wi-fi, and TCP/IP. Component standards are useful, not only in telecoms, but also much more broadly: for example, standards for headphone jacks, power receptacles and tires. Problems can arise when standards become too large in scope. Imagine standards that would prescribe the entire headphone, exactly which devices could be attached to power receptacles, or entire cars or trucks. Standards with system-wide scope will, by nature, take longer to establish and react to changes in the marketplace. But carriers typically desire or, in some cases, are regulated, to comply with these standards to provide globally interoperable communications.

In contrast, OTTs can freely determine the system architecture for their services. They do not need to cooperate with their competitors, and their constraints are simply to build on top of existing common infrastructure and interoperate over the Internet.

System standards are typically driven top-down, with specs that efficiently support the use cases that were used to form them. Operators often have siloed⁵ networks and operations based on different services that follow different standards (or came from different acquisitions). "Why bother when the vendor knows the standard best?" is a common mindset. Consequently, rather than being a source of innovation, the model has become to develop standards as an industry harmonizer. Despite hiring the brightest engineering minds, the industry has been incapable of reforming SDOs to create real-world, operator-class cloud solutions.

It is not our intent to lay all the industry's woes at the feet of standards – just to show this as a significant example of how telecoms has matured and ossified, and luckily one that we can do something about.

⁵ For example, a silo for mobility, another for metro Ethernet, and yet another for Broadband. Finally, a painful effect of reliance on standards has been its impact on attitudes towards in-house technology development

There are other drivers in the mature telecoms industry that have led to compartmentalization and specialization over time. Some of these include divestment of the industry into manufacturers and operators, separation of regulated and non-regulated services, freezing of service definitions into tariffs, network unbundling regulations, and long-term, negotiated relationships with suppliers of materials and labor. This is the backdrop to the strategic challenge and opportunity facing the industry. All these characteristics are desirable for maximizing efficiency within a status quo. However, gaining access to much more innovation, supplier options, and service options means disaggregating system standards and taking a fresh look at the business from many facets, including technology, interfaces, and component choices.

The change does not come without a cost. Adhering to system standards for networks has allowed operators to reduce their levels of engineering and design skills. However, the flipside is that operators might have also lost their technical and operational capabilities to onboard new technologies and innovate. Sadly, without this skill set, the well-intentioned efforts to advance and adopt SDN and NFV are largely without real benefit to operators. Operators are not getting the technologies they need and are unable to use the ones they get. There is not yet a proven approach or credible supplier to internalize cloud technologies, and one may never emerge if the industry follows its current trajectory.

Recognizing that the industry cannot wait, a few operators, including AT&T, Deutsche Telekom and Telefónica, are working on an alternate approach based on cloud architectural thinking and ways of working. Collectively they believe, if access networking could be rewritten as a cloud-native workload running on COTS infrastructure, it would open three significant and attractive possibilities. First, the increasing traffic demands could be better served using the more efficient cloud technologies and operations. Second, the flexibility offered by programmable, general-purpose, cloud-like infrastructure enables highly targeted elastic approaches without the need to forklift the whole system. Third, access need not be the only workload in such a deployment, and new services, from infrastructure to applications, can be made available to wholesale users and customers. The following section describes this concept and illustrates how this clean-sheet approach can enable operators to experiment with and internalize cloud technologies to positively transform their businesses.

2. The central office pod – *A winning cloud design*

To remain relevant, the industry must adopt data center technologies and align its operating model with cloud ways of working:

- The CO pod is a cloud platform, built on cloud technologies, and enables cloud ways of working.
- The CO pod uses the same open-source tools, cloud networking, data center technologies, and service mind-set that are used by cloud behemoths, as well as the same DevOps techniques to automate workload and infrastructure management.
- The CO pod design can support multiple mobile and fixed workloads, including edge computing.
- Aligning their technical platforms with hyper-scale architectures allows operators to better address data and network growth, as well as increase the available labor pool.


The central office pod, or CO pod, is a new design that takes a different approach to operator transformation. It avoids boiling the ocean with long-lead-time, multi-domain-operator transformation programs. Instead, it focuses on reimagining the operator production platform as a cloud-services platform, starting with the access network. Through clever rearchitecting of industry-standard designs and using CUPS⁶ principles, proprietary access equipment is disaggregated and morphed into a small-scale infrastructure cloud platform with specialized access “peripherals,” which provide an “access-as-a-service” software platform (see inset: “Origins of the CO pod”). However, this design is critically different from its proprietary equivalent in three ways. First, the CO pod can run access **as well as other application workloads** on the same hardware and software stack. This requires a re-think of networking as something that can be dealt with in an IT way. Second, it uses the same open-source tools, cloud networking, data center technologies, and service mind-set that are used by cloud behemoths. Third, it uses the same DevOps techniques to automate workload and infrastructure management. The combined effect of these architectural changes is creation of altogether new possibilities for operator production platform transformation. Access networks and associated IT systems represent a large proportion of industry capital spending, so even small improvements in competitiveness from one operator to the next can boost operator value generation capability and ROI significantly. The CO pod also de-risks potential cloud services and network edge services, because it does not isolate capital investment to only one purpose. Because there is a constant need to invest in network upgrades and expansions, and since there are regular technology advances that increase

the data rates in networks, there is a safe business plan to deploy cloud infrastructure solely to support network growth. Additional services on that same cloud are at little or no risk of stranding capital and can be opportunistically explored. In our collective view, because the CO pod provides a safe place that lowers the costs and risks of experimentation and learning cloud concepts in the operator production environment, it is uniquely able to drive service innovation – and therefore, it is a **winning design**.

Origins of the central office pod

The CO pod design is inspired by Open Network Foundation's (ONF's) CORD program, and also by the Open Compute Project (OCP) efforts to drive web-scale technologies into other data centers and adjacent industries, such as telecoms. Mimicking the cloud architectural pattern, it sees central-office infrastructure engineered, provisioned, and orchestrated, just like with web-scale data centers. In this new design, modules of network, compute, storage, and applications work together to deliver networking services in a repeatable design pattern that supports both fixed and mobile applications. It is a single, multipurpose architecture that allows workload pooling, which lowers cost and complexity. It has also been shown to be more capital, energy, and labor efficient. As a result, it has a track record of elastic scalability to handle large numbers of devices and traffic compared with the ability of traditional network and data center platforms. Finally, it has been shown to enable rapid innovation in one of the most competitive marketplaces: the World Wide Web.

6 CUPS refers to Control and User Plane Separation

Figure 2: Illustration of the multi-purpose converged CO pod

The physical basis of the new design is a modular, multipurpose infrastructure pod engineered for the central-office environment, as shown in Figure 2. The CO pod is not based on one-off hardware designs from telecoms industry vendors or OEMs; rather, it is based on general-purpose OCP hardware specifications, which are supplied by many vendors and used across many industries – with some diligence to ensure they can work in central-office environmental conditions. It consists of a composable rack of compute, storage, high-speed, programmable switching fabric, as well as special-purpose devices to enable FTTx access, called disaggregated OLTs; all of which are supplied by cloud industry vendors or ODMs. Like many cloud systems, there is no need for a complete consensus on several important elements, such as deployment topology and the appropriate software environment. There is room for system differentiation, even though these are constructed from similar or even identical components. The ecosystem enjoys multiple options and approaches, which enables higher levels of technology control. These include both open-source/spec and proprietary components, as well as commercial support options, with varying degrees of system integration provided in-house or by independent suppliers.

Workload convergence

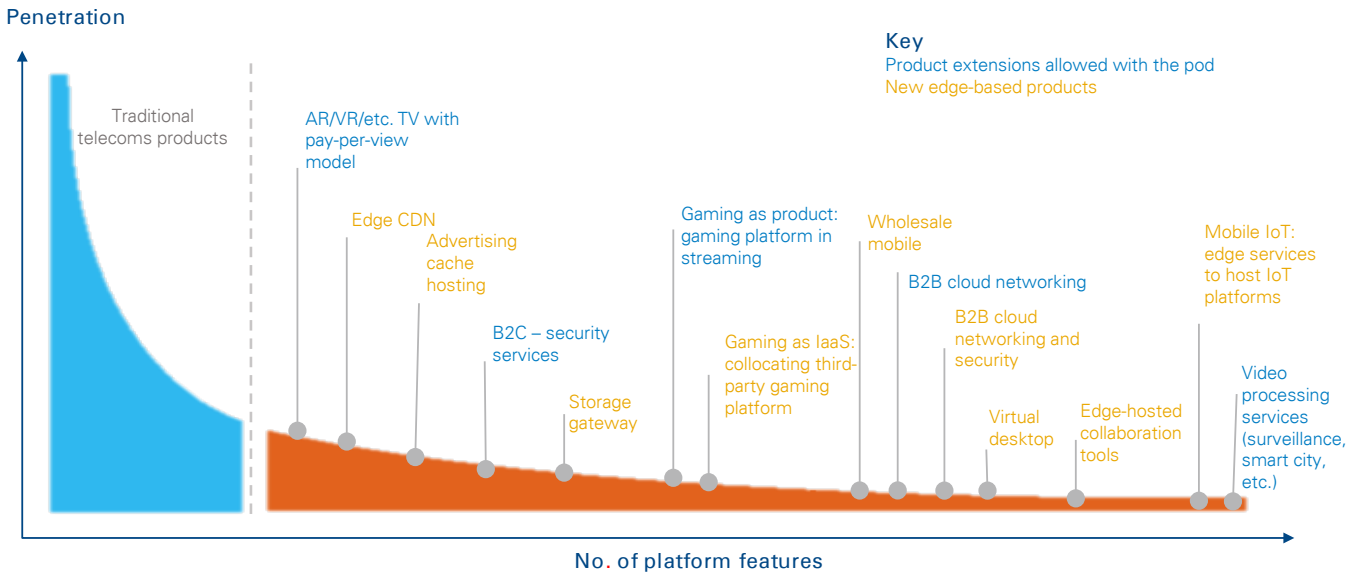
An appropriately equipped CO pod can support multiple workloads. The CO pod can focus (be economically justified) on fixed-network access in FTTx environments, and also serve as an edge production platform for smart home or office and perimeter security services. The CO pod can support stand-alone edge services, future 5G services, and edge cloud applications. Bringing together mobile and fixed workloads in a common pod allows truly converged services, enabling transparent, access-agnostic traffic aggregation and management. The CO pod can also be used for much more than simply hosting operator edge functions, including cloud

value pools. With the right security and isolation between internal and third-party workloads, spare capacity can be made available to third-party developers. Leveraging the locational advantage of the pod to provide low-latency infrastructure services, developers can deploy latency-sensitive workloads such as augmented reality and localized, data-intensive workload processing as a precursor to ultra-reliable low-latency services in 5G.

Edge computing services

Redesigning access networks also provides the tools for operators to innovate services and user experiences - emulating cloud players. Just like a public-cloud DC, the pod is a delimited infrastructure resource, which means it can be managed, provisioned, orchestrated and patched in isolation from the rest of the production platform. As a result, it provides a safe place to experiment with new product ideas and software prior to widespread deployment.

Figure 3: Emerging edge-computing workloads



New services might use the CO pod’s locational advantage as a feature to meet the needs of latency-sensitive, massive-edge data or cyber-type workloads, as shown in Figure 3. Alternatively, the ability to locally host services can be used to delocalize centralized service delivery platforms, as well as customer and order management systems⁷. This can enable a new type of quasi-autonomous production model that contains locally hosted application services; subscribers consume most of their services from central-office infrastructure. This new approach to production can provide a transient or permanent solution to the complexities of dealing with dozens of legacy services and platforms. In addition, under the right conditions, local hosting and other local micro-services can be extended to third parties, based on an open-edge infrastructure services model akin to the public cloud, to cement partnerships and capture additional revenues.

Adopting the CO pod allows operators to not only address the forces of change, but also recast themselves as infrastructure-based service companies built for service differentiation, or ultra-lean, low-cost connectivity providers to edge-based applications. Specifically:

- Aligning operator technical platforms with hyper-scale architectures allows operators to better address data and network growth. It also ameliorates changes in technology, suppliers, and generations of equipment and reduces the time to market for these types of changes. The CO pod comes along with disaggregation of vertical equipment architectures and remapping of the resultant capabilities to appropriate cloud-native micro-services, merchant silicon, and cloud infrastructure so the system can be largely

supported using typical data center equipment. This helps to bring much-needed competition into the industry supply base. Moreover, softwarization of access networks enables the use of cloud-hardened, open-source software to facilitate continuous innovation, closed-loop automation and services mashups (see inset: “Retooling for automation”). The outcome is a dramatically wider universe of hardware, as well as commercial and open-source software solutions that bring operators closer to the economies of scale enjoyed by web-scale providers.

- Adopting cloud technologies widens the addressable labor pool from which operators can source talent. The cloud paradigm is developing a vast new pool of talent from cloud and security architects, big data engineers, and agile and DevOps specialists. Their tools of choice and skills are vastly different from technical skills previously found at operators. The wider virtualization movement, which caught the industry unprepared as well as under-skilled, provides important lessons. Rather than taking an insular view to talent, the new design recognizes that the industry must pivot towards the same platforms and tooling that are common in the cloud; this will enable the industry to draw on this wider labor pool, which will bring with it vital skills and ways of working to drive individual as well as industry competitiveness.

⁷ Also referred to a BSS “Business Support Systems” including customer portal and apps

Bringing cloud technologies and practices inside the telecoms operator allows the industry to align itself with the cloud paradigm. It is not yet apparent whether operators should replicate existing cloud models for innovation and monetization or find other ways. Whatever the direction, the architecture enables operators to fast-track transformation as well as have a go at innovation, creating meaningful differentiation among operators. However, for the CO pod to become reality, it must be industrialized, productized and deployed. In the next section, we make the strategic and economic case to invest in doing just that.

Retooling for automation

It's critical to drive automation into the design, deployment, and operations lifecycle of the CO pod. Looking at Figure 3 and taking penetration as a factor in determining revenue, it should become clear that developing bespoke approaches to the most popular of services may be possible. However, that approach adds increasing overhead as you move along the long tail and will prevent deploying profitable services at some point.

When developing, using, and reusing common automation across services, the incremental expense needed to support the next service is only in its use of infrastructure resources, not any additional design, deployment and operations cost. In short, the ability to support and monetize long-tail workloads is contingent on the ability to drive end-to-end automation.

3. The case for the central office pod

The CO pod provides operators with a safe place to virtualize/re-engineer existing services, as well as prototype and test new service ideas, using cloud-hardened development and operations methodologies. The CO pod gives operators a safe place to start over; we see three options:

- “Converged virtualized access” can encompass both fixed and mobile access, harmonizes the way all traffic is treated at the edge, and provides higher throughput at lower total cost of ownership.
- “Autonomous operator” opens a range of options to deploy complementary, highly-automated edge services using the CO pod.
- “Open operator platform” enables qualified third parties to exploit the edge using pay-per-use models in a model akin to public cloud.


No two operators have the same priorities or starting point, so a thorough economic analysis must consider the specific competition, technology roadmap and debt for each case individually. Nonetheless, to illustrate the logic and benefits of onboarding the CO pod, we describe *three pathways* for product and production platform development, as shown in Figure 4. The economic logic of each pathway is determined using a side-by-side comparison of proprietary equipment designs with comparable functionality that were built using disaggregated and virtualized data center equipment (i.e., procurement levers only). However, the strategic benefits of the CO pod architecture require a more holistic view of the changes enabled by the new design, such as agility and efficiency in developing new offerings, channels, operations and supply chain value. These changes provide the foundations to drive step-change improvements far beyond procurement, including end-to-end production platform efficiency and innovating the

operator business model. Nonetheless, such direct comparisons are valuable because the latter is subjective.

Converged virtualized access

The first approach, also a segue into the other two options, is to use the CO pod for virtualized access as defined today within the existing technology landscape. Virtualized access can encompass both fixed and mobile access and harmonizes the way all traffic is treated at the edge. This saves backhaul cost, reduces interfaces towards the backbone, and enables hybrid access at almost no additional cost. The economic logic behind this choice is that the pod provides higher throughput at lower total cost of ownership compared to classical industry equipment designs. At the same time, it simplifies the process of integrating the technology into an operator network. Because the CO pod is based on commonly available, open hardware and

Figure 4: Option space enabled by the central office pod



software, there is a great amount of transparency into multiple areas of concern. Hardware BoMs are known and available from multiple suppliers, so design-to-cost, fair pricing, and vendor lock-in can all be managed effectively. Software is also open, so the operator is much more able to manage similar concerns about the value being provided by a software integrator or supplier and, very importantly, vendor lock-in. This first approach has arguably the lowest technical risk of any of the approaches but limits the rewards to efficiency gains for providing broadband access and access convergence. For many, this is a no-regrets rationale to get started on deploying the CO pod, with the next two approaches becoming bonus benefits that can be pursued in the future with little additional investment.

Autonomous operator

As an extension of the virtualized access concept, the *second approach* utilizes the full set of technical capabilities of the CO pod, expanding the scope to include non-access operator services. The rationale behind the choice is exploiting the platform capabilities of the pod to shift other telco services towards the edge. This is not simple; it requires taking a hard look at existing processes and systems to determine a systematic approach to service and operations transformation. Key to this effort is the creation of marketable, differentiated and/or low-cost edge services, while using the change to reduce dependency on legacy platforms. Operator investment pivots from simply providing bandwidth over long distances to supplying infrastructure that allows applications to consume less bandwidth because there is less distance to their users.

It is important to note the potential benefit of distributing services, in that it can help simplify OSSs, as well as service logic and design complexity (see inset: “New options for distribution”). Many of today’s services are supported with sets of network elements distributed across several offices – like beads on a string. Such services need to be managed as a distributed system, and that pushes a lot of the management and service assurance into the central OSSs that oversee multiple locations. Moreover, fault tolerance often requires geo-redundancy and service state to be duplicated in multiple locations. When a service is largely or completely located at the edge, distributed service assurance and provisioning becomes a local matter and the OSS is simplified at the central tier. This distributes complexity in a more manageable way. Moreover, as for geo-redundancy, in most cases there is little that can be done when the serving office has a systemic failure, since the customers are not connected to any other office, and therefore there is little need to consider redundancy beyond that locality. Note that this argument considers that there is still a highly available WAN IP network attaching to these edges.

New options for distribution

For decades, common sense for operators has been to centralize when they can and distribute when they must. The second approach challenges that logic. Rationale for centralization has been based on aggregating service workloads into larger, more efficient, less fragmented network elements, and it still holds true for legacy box deployments. However, with workloads becoming software in computers or slices projected into merchant silicon switches, it has become possible to distribute fine-grained amounts of service logic without fragmenting the capacity of edge cloud infrastructure. In this environment, a new preferred topology emerges. That topology distributes service capability and logic to the point it’s delivered, and that point is at the edge of the network. However, databases for service entitlements, data lakes to feed automation engines, and principal management and control operations facilities remain central.

Open-operator platform

An independent *third approach* is to use the CO pod for edge cloud services. This is another logical extension of the first approach. The economics underpinning this choice is the ability to sell excess infrastructure capacity for use by applications that benefit from being local (see inset: “Centralized versus edge cloud workloads”). “Open platform” logic enables qualified third parties to exploit the edge infrastructure using pay-per-use models that are commonplace in the public cloud. This could result in new forms of collaboration with third parties and/or incremental revenues from either rental of the infrastructure or provision of data services. The approach is presented as an add-on because it is seen as having significantly more business risk than a stand-alone approach. On its own it becomes a “build it and they will come” proposition. However, as an add-on to one of the previous approaches, the additional business risk is largely eliminated.

Centralized versus edge cloud workloads

There is not a single broadly accepted definition for edge cloud, so defining how it fits into the existing cloud infrastructure landscape is challenging.

There is also still a lot of uncertainty around the potential benefits of edge-cloud use cases. In particular, the question of how far to the “edge” the compute and storage resources need to be placed is highly controversial and depends on both the workload and the specific deployment situation, e.g., the geographical extent of the network or backhaul delays. While in some scenarios CO pod co-located data centers yield net improvements in user experience, in other scenarios a few centrally located data centers per operator and country would suffice to achieve the same results. For this paper, however, we define the edge cloud as existing in some fraction of serving offices – offices where access technology must be placed. This is for two reasons: 1) it allows sharing capital investment with access, and 2) it reduces the risk of competing against centralized cloud offers from web-scale providers.

The choice of either, or both, of the latter approaches will be driven by factors that will vary from one operator to the next. For example, if an operator just rolled out a new service in the traditional way, it might want to wait for that service’s lifetime to come to an end before moving it to the CO pod.

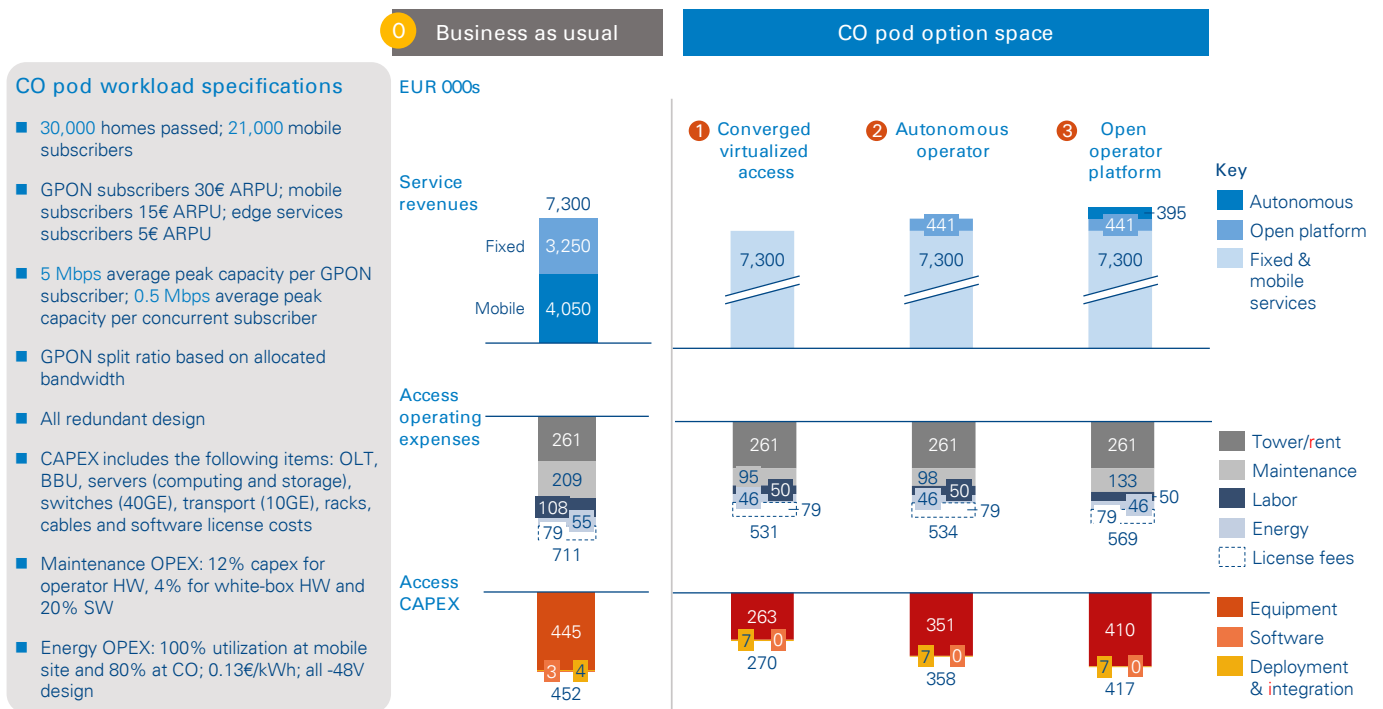
Similarly, introducing third-party workloads to the CO pod requires operationalizing the infrastructure as a public cloud, with security, billing, and customer management tools that go beyond what’s needed for a private cloud.

Economic impact

In order to provide a first-order approximation to the sorts of economics associated with the CO pod and the approaches that have just been described, Arthur D. Little has developed a model that compares these approaches based on public information and an exemplar CO pod described in the technical section of this paper. The model is used for all the economic claims in this paper.

Figure 5 illustrates the link between the different approaches and the addressable value levers, comparing industry-standard design with the corresponding CO pod design. These benefits are contingent on service providers succeeding in creating a viable ecosystem of suppliers that support the technology. The scope of the cost analysis covers network aggregation, subscriber management, and mobile baseband processing, as well as mobile- and fixed-routing functions. The visualization is based on a simple model to illustrate the delta change impact of revenue and cost levers associated with each option. It also illustrates that access platforms represent a small proportion of revenues. The economic case has been built based on bottom-up cost comparison of traditional equipment and virtualized access platform economics, using representative industry costs in a greenfield scenario.

Figure 5: Economic impact of central office pod deployment options



Source: Arthur D. Little estimates and calculations

Figure 5 also shows that the impact of virtualized access is mainly operating and capital efficiency. Putting things into the overall context shows that the aggregate results are small and highly dependent on automation benefits derivable from changes in operating practices. As detailed in Figure 6, according to the Arthur D. Little model, the CO pod has circa 40 percent cost advantage in CAPEX, but 25 percent in OPEX. The former is due to its reliance on architectures based on open-source hardware and software, while the latter illustrates the relative importance of central-office non-equipment costs. The largest CAPEX differences are in the relative cost of fixed traffic aggregation and routing, as well as GPON optical equipment (new vs. old cost).

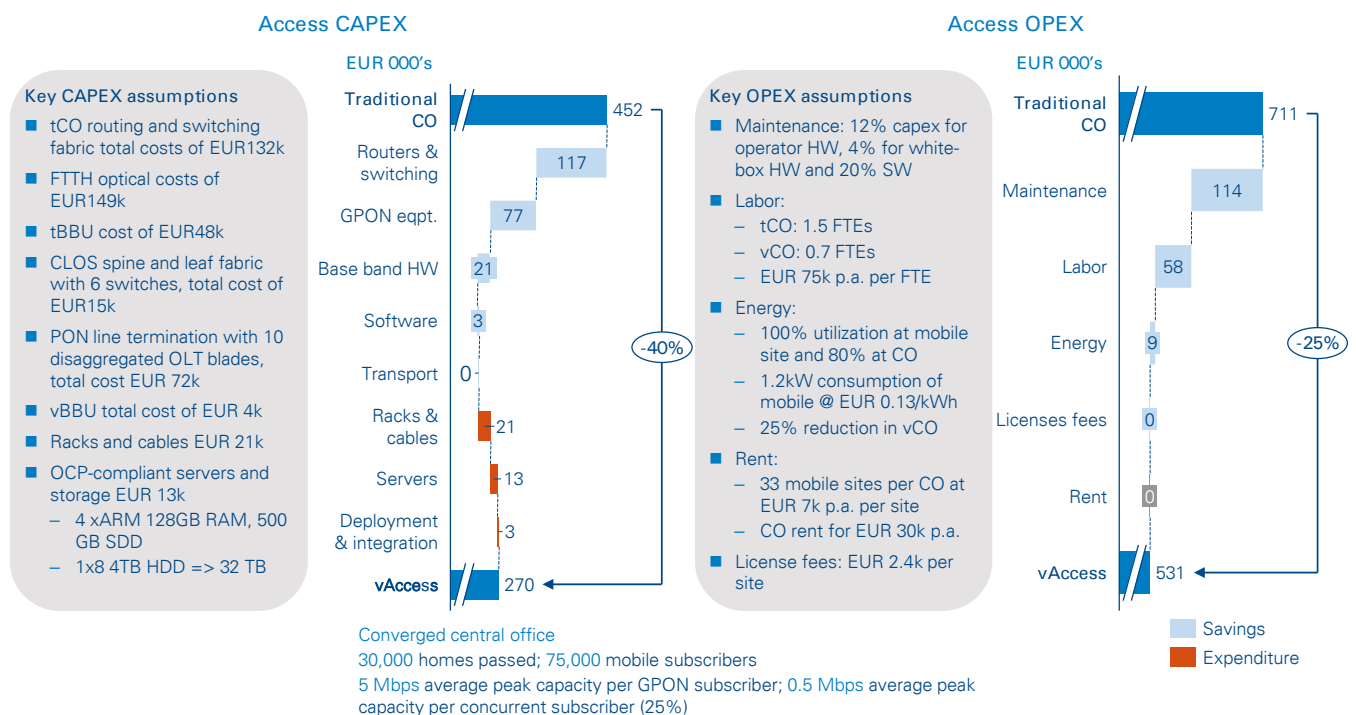
The analysis is based on published price lists and shows that white-box, programmable switches cost one-tenth as much as typical operator routing boxes, whereas commodity optical equipment costs are less than one-half. The model also contemplates virtualizing radio-access processing for mobile networks. Though still emergent, an estimate of marginal cost reductions of 80+ percent is enabled by pooling hardware resources with the pod control functions. Major OPEX differences are more indirect. A large proportion of the savings are based on lower maintenance costs associated with commodity hardware. Further savings are based on greater automation and energy savings. It is important to note that the CO pod does not support all telecoms control functions and protocols; rather, it is a minimalist, software-based design that provides only what is necessary to deliver with carrier-grade reliability.

“Autonomous operator” and “open platform” are revenue plays that use marginal economics to unlock incremental value from the platform. While edge platform services relevant to third parties are a work in progress, numerous spaces linked to video and programmatic networking capabilities are emerging. Figure 7 illustrates numerous sources of incremental revenue gains that could be possible using proximity and low latency. Along both pathways, we expect similar costs to access virtualization; however, we expect higher integration costs associated with the complexity of creating new operations and services platforms, distributed over multiple pods. The upshot is that the CO pod can be used to gain incremental share of wallet, as well as market share. We have attempted to quantify the former, envisaging up to 11 percent gains over and above connectivity revenues based on proxy to comparable cloud services. While these gains will require 5+ years to achieve, they illustrate that such incremental gains could be used to partially self-finance the shift to a CO pod architecture.

The precise economics for incumbent and challenger operators will be different. Incumbents are endowed with greater density of central offices. Challengers, in contrast, rely on regulatory access regimes and new build-outs, so the economic drivers for change will vary. However, this creates new options where challengers may not necessarily be the underdog.

Most incumbents will find they have significant spare capacity in existing central offices; however, significant capital investments may be required to expand or upgrade facilities. On the other hand, challengers will find it prohibitively expensive to match


Figure 6: Relative economics of the CO pod



Source: Arthur D. Little estimates and calculations

Figure 7: Illustrative edge cloud value pools

EUR 000's



Source: Arthur D. Little estimates and calculations

incumbent deployments site-by-site. A more suitable alternative would be collocating all or part of their infrastructure into nearby operator-neutral or operator-peering facilities. In addition to enabling them to access competitively priced, DC-grade space, it puts challengers in prime position for direct peering with tier 1 web-service players, as well as deploying unmodified, high-density standard racks for the CO pod. Moreover, the ability to be a magnet for peering partners could be used to secure ultra-competitive collocation services. As video becomes the dominant form of internet traffic, the impact on transport costs could be substantial.

The new architecture creates new efficiency and revenue opportunities but is contingent on the CO pod becoming a de facto standard. It's especially important to the third pathway, in which third parties make use of edge-cloud capabilities. If few operators support this capability, it will be less attractive to third parties for developing for the edge cloud. Given this chicken-and-egg situation for edge-cloud demand, it makes sense not to invest in edge-cloud infrastructure without de-risking it by placing the virtual access workload on the same infrastructure. However, getting the CO pod into production will require letting go of industry handrails and accepting new, unfamiliar risks. Additionally, unlearning old habits is hard, and will require leadership that must come, not from technology, but from business decision-makers, as we describe in the next section.

4. Leading the creation of a *carrier-cloud* mind-set

Operators must decide if they will lead or follow. For “those who dare,” the prize could be significant:

- “Followership” means waiting for a comprehensive alternative supply base to emerge at some unknown future point in time.
- “Leadership” means taking matters into one’s own hands: technology must master engineering and operating cloud infrastructure to provide telco-grade services; sales-marketing-product must evolve in order to exploit capabilities to develop new, exciting products and evolve supplier ecosystems.
- Operators must create an environment that allows teams with the right skills and mindset to experiment (without fear of failure).

Making CO pod-based production a reality will require profound changes for operators. Operator technology organizations must master the details of engineering and operating cloud infrastructure to provide telco-grade services. In parallel, the sales-marketing-product function must evolve in order to exploit the unique capabilities to develop new, exciting products and future platforms for growth. All this is easier said than done because there is no proven roadmap, standard, or large vendor available to help.

The scale of the change associated with access transformation is not unique. Other industries have seen arguably similar shifts that challenge the status quo: the introduction of renewable power into a traditional utility value chain, that of the electric powertrain into conventional vehicles, and the ongoing introduction of blockchain technologies into banking. Therefore, operators should not expect the path to transformation to be smooth or uneventful. However, once the dust settles, the business pay-off can be substantially lower costs, extensions and changes to the business model, and even globalized operations or a “wow” customer experience. Lessons learned include the importance of creating space or labs for experimentation and learning, as well as that of creating oversight mechanisms that encourage technology adoption to solve high-value business challenges.

Adopting the CO pod architecture requires making decisions and accepting their associated risks. The pod is a pre-production technology with a wide but shallow vendor ecosystem to support its development. Therefore, each operator must decide whether it is willing to wait for a specialist supply base to emerge or take matters into its own hands by helping build

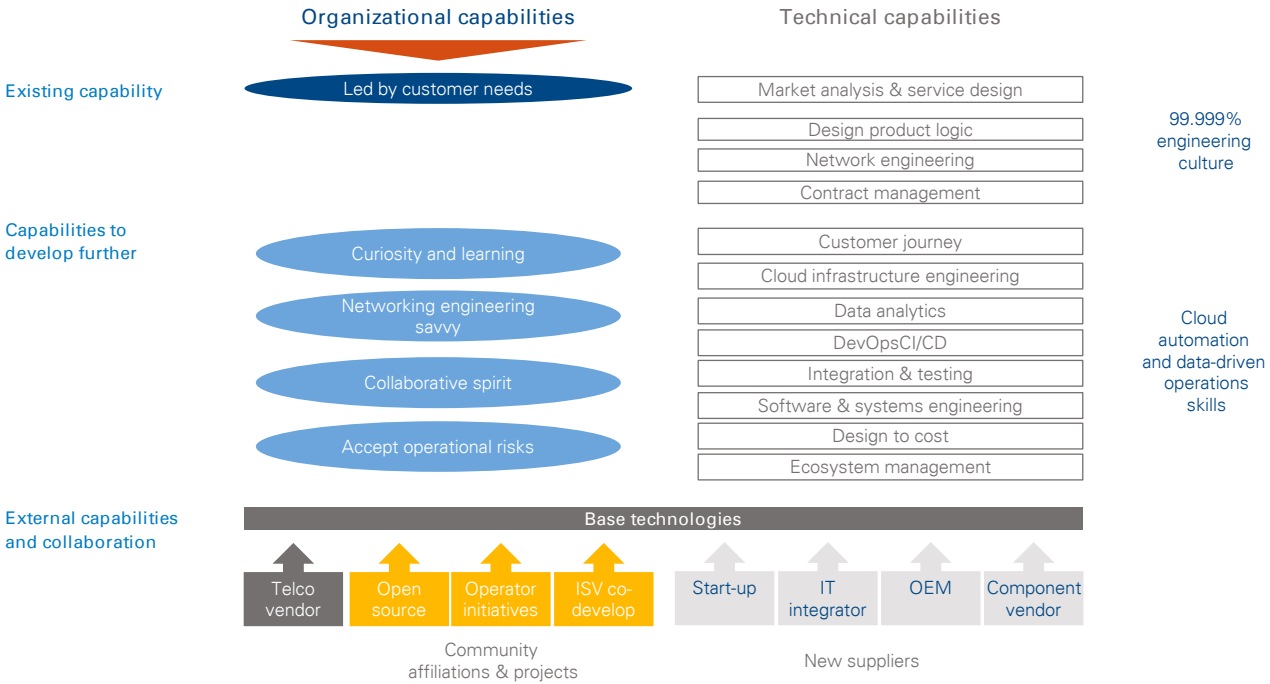
the platform. Followership may result in missing the boat, but taking control means assuming greater responsibility for the end-to-end engineering. It is not a decision that technologists can make alone; rather, it is a decision regarding operator corporate strategy. Variations such as scope of activities, market position, regulatory constraints and level of control over existing operations will determine the right approach. The operator organization must evolve and redefine what strategists, product marketers, technical managers and operations personnel will do differently, and put structures in place to execute. This requires rebuilding the product development model, as well as the operator technology supply base.

Success requires operators to structure teams with the right skills and mind-set, as well as create a supportive ecosystem of partners and vendors, as shown in Figure 8. This might even require them to reacquire skills lost in the last decades ([see inset: “Technology capabilities”](#)).

Technology capabilities

Until the mid-90s, leading operators would often play a key role in the engineering and development of technologies used in their networks. The famous 5ESS electronic switching system, originally developed by AT&T, is a powerful example. Other examples from the co-authors include Deutsche Telekom’s development of ISDN and Telefónica’s co-development with ITT (acquired by Alcatel) of the S12 switch. However, for several reasons the eruption of IP technology has seen operators’ role diminish, with companies relying almost entirely on a handful of vendors for their core technology needs.

Figure 8: Capability and ecosystem map



Taking advantage of the new platform will require savvy product marketers and service engineering to seize the unique advantages of the new design and mitigate new pitfalls and risks. Operators with narrow focus on access transformation must develop and exploit the new platform within existing network and IT operational environments by developing capabilities to wrapper the platform in order to benefit from the automation tools with the least integration effort for the existing OSS. Operators that see virtualization of the central office as an autonomous edge production or open production platform must rethink their entire operational systems and billing architecture to enable rapid experimentation focused on delighting customers. This will require empowering the CIO, as well as product and network organizations, to find solutions, and asking for results. De-commoditization of operator products and services will require learning to learn again. Rebuilding the product function and platforms, developing cloud architecture engineering, and learning to engage with customers all require patience. All approaches require learning to attract external talent, as well as reactivating internal staff. Attracting external talent will require persistence, a credible story and a space to work. On the other hand, reactivating staff will require dealing head-on with cultural issues to encourage self-development and taking professional risk. Few have attempted to reskill the entire workforce like AT&T, and the company provides a template to emulate (see inset: "Organizational skill upgrading").

Organizational skill upgrading

Executing in the new landscape requires skills in cloud-based computing, coding, data science, and other technical capabilities – but talent is limited, and everybody is going after it. AT&T is refocusing employee education and professional development on reskilling technical staff. In collaboration with Georgia Tech and numerous online course providers, AT&T is providing opportunities for almost all staff to acquire the skills they and the company will need for the future. It is not optional, because, in effect, every employee is being asked to requalify themselves for the jobs that will be available in the next decade. This is just one element of an overall talent, capability, process and culture change that AT&T has undertaken.

For the new design to become a reality, supplier ecosystems must also be regenerated. Unbundling the technology stack means companies must figure out what components are needed to build, deploy and manage operator-grade virtualized access networks – and how to source them. In this future, operators may not get whole-hearted support from all the traditional vendors, which will be understandably reluctant to erode profit pools and disturb the status quo. This means looking toward the wider cloud landscape with willingness to innovate models of sourcing and collaboration key.

Regenerating the supplier ecosystem will require trial and error rather than an ex-ante decision. It will be an evolutionary process requiring an understanding of existing commercial as well as open-source solutions. Elements will need to be developed by each operator to ensure technology control and differentiation. This will require “learning by doing” and allowing natural selection of players which are trustable and reliable and have staying power in their niches. Operators will need to build the capabilities to shape and orchestrate enhanced sourcing ecosystems, including those developed by cloud service providers.

Cloud hardware and software suppliers are very different from telco suppliers. They are typically smaller, specialized system suppliers whose responsibility is limited to the component. Their limited scope of activities means they are unlikely to be able to match the breadth of services the operators are accustomed to when working with established vendors. This inevitably changes the split of labor and responsibility between in-house teams and third-party vendors. Success requires resolving these issues and managing the diversity to produce a homogenous service. This should not be foreign to operators with long memories – as illustrated in the “Technology capabilities” inset.

Each operator does not have to go its own way. Recognizing the change is a common challenge facing the industry. Operators can pool technical and financial resources and direct them toward creating one or more ecosystems. In these new ecosystems, operators define the rules of engagement, ensure the industry gets what it needs, and create opportunities for others to build new businesses and business models. This is the essence of a community-based approach: a number of operators working collaboratively on a shared ambition, focusing on common elements that do not drive differentiation. In this community effort, success will create the conditions so these common services can be consumed or acquired from third parties.

An issue that must also be dealt with is the role of traditional telecom equipment vendors. What should operators expect from them in this new future? These vendors are the backbone of the industry, and finding new win-wins is important to sustaining existing telecoms’ asset base. Regenerating the existing telco supplier ecosystem will require far more than standards-based cooperativism. It will require technology openness and collaboration to identify new sustainable development options⁸. Such options could include prime integrator, as well as co-development, open-source software

support and evolution of legacy equipment to support emerging open-source, de facto standards. The latter requires an entrepreneurial, not a standards mindset on both sides. Given the steep learning curve associated with skills and ecosystem development, we think “wait and see” is not an option. The impact stretches far beyond the network into product features, customer support and the business models used in the industry.

The impact of the new architecture is far-reaching. It requires a concerted effort by decision-makers to create and sustain the conditions for success. Moreover, the new design must necessarily co-exist with *business as usual*, and they must complement each other in direction, capability and ecosystem development. Creating space for both to coexist will require leadership and cunning to drive eventual convergence. In the next section we share how AT&T, Deutsche Telekom and Telefónica have approached the journey.

⁸ One example of just such collaboration is embodied in the O-RAN Alliance. Operators lead and define the new model, but with co-operation and co-creation from existing suppliers in the RAN market

5. Launching a *minimum viable* access transformation program

Strategic priorities, not technology, must drive access-driven transformation programs:

- There is not a single overarching approach to framing an access transformation program.
- AT&T wants to systematically embed cloud technologies in all areas of the production platform.
- Deutsche Telekom is focused on bringing down costs of future access platforms, as well as broadening the supplier ecosystem.
- Telefónica is concentrating on using the platform’s edge-cloud features to create new services and revenue streams.

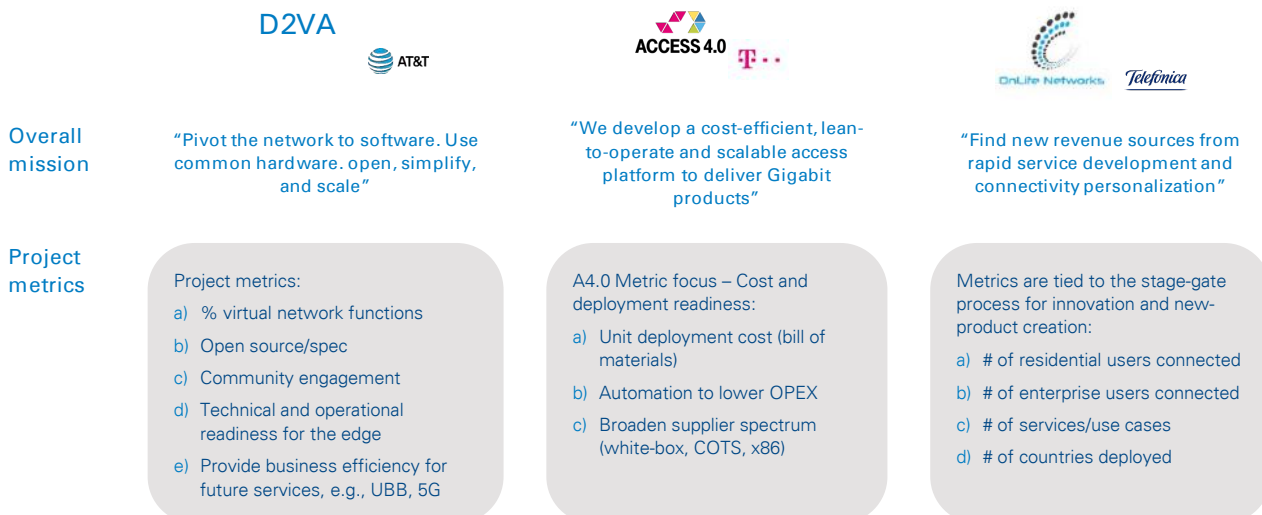
Creating space for radical innovation and breakthrough solutions in traditional environments is not trivial. To encourage community and accelerate industry-wide adoption, this section provides unique insight into the strategic context of the AT&T, Deutsche Telekom and Telefónica programs and how they have approached program design. Their journeys are still being undertaken, but they provide important insight into problem framing and monitoring, as well as guidance on a number of other topics that are important for execution.

concert with the global program to share reusable assets. In contrast, Deutsche Telekom’s approach is focused on bringing down costs of future access platforms, as well as broadening the supplier ecosystem. It sees the cloud as a tool, and not necessarily an end state. Finally, Telefónica is focused on using the platform’s edge-cloud features to create new services and revenue streams. There is no one-size-fits-all approach. Focus is determined by perception of value, market position, network deployment windows and technology prowess.

Figure 9 summarizes the differing priorities for access virtualization at AT&T, Deutsche Telekom and Telefónica. In a nutshell, the AT&T approach is part of a larger program focused on systematically embedding cloud technologies in all areas of the business. This means the access team works in

The remainder of this section describes how AT&T, Deutsche Telekom and Telefónica have structured their minimum viable access-based transformation programs.

Figure 9: Highlights of AT&T, Deutsche Telekom and Telefónica access virtualization programs





In response to an internal business and technology roadmap program, AT&T determined that there would be profound changes in the wider competitive environment. In response, the company's senior leadership launched a network and IT transformation program that focused on customer self-service and speeding up time to market. Led and funded by the Technology and Operations organization under the heading of "Domain 2.0" ("D2.0"), the transformational initiative was created to *pivot the network to software* across all network domains and associated operations. With direct oversight from the officers, D2.0 went public in November 2013 when the D2.0 vision white paper⁹ was released. The paper describes AT&T's vision for its future architecture and mode of operation. It triggered several sizable investment programs focused on comprehensive retooling of many aspects of the business. Three business imperatives are pursued by D2.0: Open, Simplify and Scale.

The initial focus was two-fold: to transform a significant portion of networking equipment, typically appliances, into software virtual network functions (VNFs) that would run on data center servers; and to establish a large-scale system automation platform that would plan, deploy, monitor and coordinate functions among many physical locations. This helped to automate and virtualize the existing production model. To advance quickly, RFPs were augmented with short, simple, yet broad market surveys – sent not only to legacy and holistic vendors, but also to start-ups and small suppliers that could only provide components of the overall solutions.

The process is called agile engagement because, like agile development, it follows the pattern of taking small, iterative steps toward a goal. Rather than relying solely on a large waterfall RFP and then instigating only a single interaction with respondents, the agile engagement process allows interacting with many more potential suppliers, and then develops potential solutions through more specific and focused interactions. It also whittles down the field of potential contributors toward a solution. By December 2014, these arrangements were formalized under the D2.0 vendor program, and an ambitious plan was communicated to staff and the outside world for network functions virtualization. Finally, in one of the largest-ever corporate retraining programs, AT&T started to re-educate its entire technical community with critical skills needed for the future. Employees were, and still are, rigorously trained to meet the needs of new job positions through leadership courses, degrees, and "nanodegrees." To date, D2.0 has resulted in investments worth several billions of dollars in SDN and NFV, as well as numerous large-scale software projects to manage the new landscape, including ECOMP (open sourced as ONAP). This

has enabled automated, end-to-end management, and more recently, DANOS, Airship and Akraino.

While initially focused on appliances, D2.0 also seeks to open, simplify and scale other network capabilities. Often, this requires breaking apart existing boxes so that they can be disaggregated into their simpler, more reusable subcomponents. Domain 2.0 Virtual Access ("D2VA") is the access component of the D2.0 program, focused on disaggregation and virtualization of the access network. Access disaggregation was considered a hard problem, so it was not initially a priority; instead, AT&T chose to innovate with an external team from ON.Lab. Based on the learnings from CORD trails, AT&T launched D2VA in 2017 and chose to collaborate with other carriers to create a production-ready design called SDN-Enabled Broadband Access (SEBA). SEBA significantly enhances operations capabilities, interoperates with legacy systems, and provides a path to an all-orchestrated future architecture. D2VA is funded and managed by the wireline and wireless architecture and planning functions at AT&T. The program is organized along three workstreams: infrastructure, wireline and mobile access. The common infrastructure workstream is aligned with other D2.0 infrastructure programs and benefits substantially from existing investments. The wireline stream is where the bulk of staffing is being concentrated. The relative scale of the wireline team is a result of the solution having matured enough to move from PoCs/trials towards deployment. The third work stream, focused on mobility, is currently in trial phase, with the goal of converging with the wireline work as it matures. Most of the team members are existing and retrained staff from within AT&T, complemented by a few externals, college hires and interns to fill specific developer gaps. The success of D2VA is tracked along multiple dimensions, including aspects such as the percentage of open software and the use of white-box hardware based on OCP specifications. While the specific D2VA metrics are focused on relative economics of deployment and operations compared with traditional solutions, the program also builds technical and operational readiness to support future roll-out programs, such as 5G, and builds an ecosystem with community engagement and support. Because of the latter, AT&T spends considerable resources to foster community relationships. The D2VA team closely collaborates with ONF/ On.Lab, OCP and the Linux Foundation (ONAP, Akraino, DANOS).

The long game for AT&T with D2VA and the larger D2.0 initiative includes attracting talent, as well as new capital, into the industry supply chain. AT&T feels that disaggregation, modularization and community collaborations will lower barriers to entry. New entrants can be adjacent providers from Enterprise IT, SI houses, traditional IT suppliers, ODMs, and even VC-funded startups. Similarly, new talent can come from IT and cloud backgrounds in addition to telecoms.

⁹ Domain 2.0 White Paper is available at https://www.att.com/Common/about_us/pdf/AT&T_Domain_2.0_Vision_White_Paper.pdf



To support the development of the gigabit society in Germany, there are numerous groups focused on disruptive thinking across Deutsche Telekom. Their areas of interest include 5G, mobile edge computing and rethinking access networks. Access 4.0 (“A4.0”) is one such program; its mandate covers fixed and mobile access networks, as well as support for mobile edge-computing infrastructure, but initial focus is on FTTH and FTTN. A4.0’s primary goal is to redesign access networks to reduce vendor lock-in, and drive “step-change” reduction in life-cycle costs through the use of data center hardware, open software and automation. The approach is based on “design to cost” principles, which are commonly employed in product development and manufacturing. It is producing a decoupled design using commodity data center hardware, software and management concepts, with the goal of supporting any existing access service transparently. Focus on transparent access adds some complexity but has allowed the A4.0 team to operate with little involvement from the product or commercial functions.

The program is hosted within the German Engineering organization with strong support from Technology & Innovation (TI) functions and teams. By design, presence within an engineering function provides the right controls and incentives to ensure direct linkage of A4.0 scope to specific business objectives, as well as alignment with Deutsche Telekom’s roadmap. It also provides controlled flexibility to pursue alternate technology pathways, procurement and staff hiring. Direct purview by an operational technology function provides controlled architectural liberty to ensure designs are production ready. Early involvement of procurement has enabled sourcing from non-traditional vendors that are more inclined towards open technologies. Locating the program in the engineering function has also simplified the process of attracting the right talent. Despite being a local program for the domestic BU of Deutsche Telekom, the technical designs are extensible to other geographies across the Deutsche Telekom footprint where applicable.

Rather than building a full set of capabilities in-house, A4.0 prefers to rely on internal staff supported by partners with deep software engineering skills, so it has been critical to select the right individuals and partners to create positive incentives to collaborate. A4.0 was able to transversally win resources from teams specialized in IP networking, OSS, operations, planning, security and software development. It is currently working closely with Reply as main partner, among others, which has been willing to collocate staff at Darmstadt and Berlin and bring in their software engineers with skills in software architecture

of large-scale systems, data science and engineering, container networking, and software-defined networking.

A4.0 is targeting production end of year 2020. There is a disciplined approach to monitor progress of three KPIs: (1) unit deployment cost; (2) automation to lower operating costs; (3) broadening of the vendor ecosystem to encourage adoption of commodity elements. Progress is monitored through systematically tracking the quality of the vendor ecosystems willing to support A4.0 roll-out, as well as continuously updating the cost-estimate model to validate direction of efforts.

Telefonica

To encourage a culture of “intrapreneurship”, Telefónica has a systematic process for new-product and -business creation under its digital organization. At least once a year, the Digital Product Innovation Team launches an “Innovation Call” to identify specific research areas of interest that could create new revenue streams from TEF core assets. All TEF staff are invited to participate by proposing innovative products, services and experiences linked to these themes.

In 2016, under the “Customer Centric Networks” theme, an ad-hoc team that combined network architecture, planning and R&D pitched OnLife™. Though OnLife™ has many possible development angles, its primary focus is on finding new revenue from rapid service development and personalization of connectivity. The project was selected for initial exploration in May 2016, after successful completion of an initial stage gate. Prototyping was approved in July 2016, and thereafter, OnLife™ went through the process of technical verification until field trials were authorized in September 2017 and the first Zona Beta clients were connected in June of 2018.

OnLife™ operates as a “virtual company”, with its own CEO, CTO, and head of architecture, as well as services and technical teams. OnLife™ execution is organized around four disciplines: Infrastructure, Edge Platform, Use Cases and Access Network. The Infrastructure team is responsible for all aspects related to access, switching and compute infrastructure, as well as the virtualization software or VIM stack. The Platform team defines and develops the APIs for internal use, as well as by third parties, and is responsible for catalogue definition, service orchestration and billing. The Use Case team is responsible for business development and supporting use-case prototyping. As of June 2018, there were 10 staff working on the Infrastructure discipline, five each on Platforms and Use Cases, and two on Mobile Access Networks. The majority were lateral hires from Telefónica. However, in contrast, most of the Use Case teams were external hires. The OnLife™ team collaborates

with other network transformation programs at Telefonica, including Internet para todos¹⁰ and Unic@¹¹, which are focused on low-cost radio access networks and virtualizing the core network, respectively. Collectively, they are trying to figure out how to advance the disaggregation and softwarization agenda at Telefonica and drive network cloudification together with extreme simplification of operational practices.

Lean Elephant™

The selection process consists of a four-minute stand-up pitch to the top 30 managers of Telefónica, who collectively decide whether to fund the project or not. If selected, the team is given three months to show the technical and economic viability of the solution through consulting with thought leaders at Telefónica and potential partners. If shown to be possible, the team is asked to list its technical and commercial hypothesis and provided with funding for prototyping. Typically, this phase lasts between six and 12 months and consists entirely of validating the key hypothesis. Subsequent steps include either beta testing and development, industrialization, and testing with real clients, which is also called the product phase and when the industrialized product is transferred to an ob. The process is a Telefónica variant on the Lean Startup™, called Lean Elephant™, and is funded centrally by the digital team after conceptualization phase. Since the launch of the innovation call, dozens of ideas have been submitted and three have reached the marketplace.

Despite its start-up structure, Onlife™ is funded 100 percent by the Telefónica group. Funding disbursements are based on progression through a formal stage-gate process. This internal proprietary process, called “Lean Elephant”¹² (see inset: “Lean Elephant”) is based on a Lean Startup™ process in which the three reference pillars of the framework applied to innovation are: (i) **Start small and aim high:** The level of ambition of the innovation projects must be high. They need to bring the possibility of global reach and the potential to make an impact in everyday life and business. This does not mean they will burn lots of resources to start with, or that they need to show full potential from day one – quite the contrary. Projects, especially at the beginning, work with bare-minimum resources, and investment increases as the project progresses, with validated learnings. The less uncertainty, the more budget. (ii) **Iterate fast to achieve efficiency in each of the maturation stages:**

This means scaling down initiatives that are too early in time, immature or unfocused, while fueling up ones that show traction. Therefore, product investment decisions conducted along the process rely not only on technological trends, but also on profound understanding of which markets digital customers will participate and spend in, in the upcoming years. Finally, (iii) **Fail fast, fail cheap and make sure you learn along the way:** Instead of devoting large quantities of energy and resources to increase the individual-success chance of a few projects, Telefonica lowers the overall risk by minimizing the failure cost for each project.

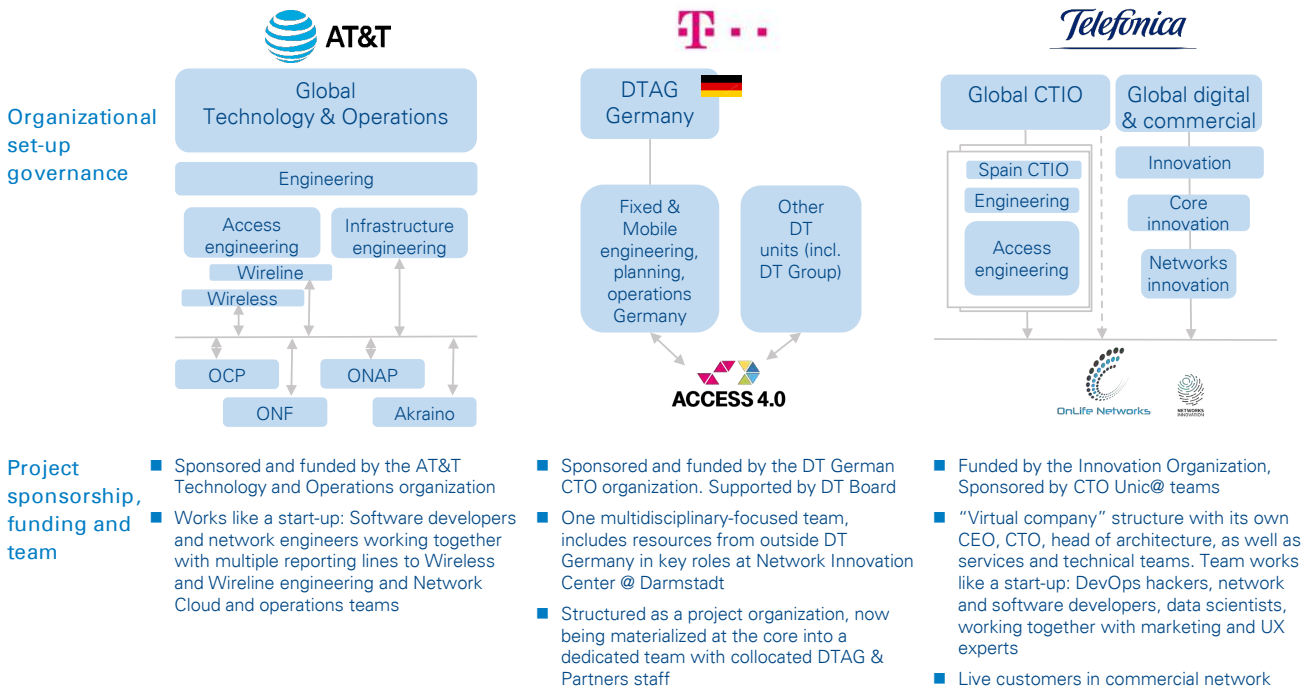
Figure 10 illustrates that there is no single-best appropriate approach to access transformation. AT&T and Deutsche Telekom have collocated the team within their operating-unit engineering organizations, with oversight from their Global CTOs. On the other hand, the Telefónica initiative is structured as a virtual company with funding from the Innovation function and sponsorship from two operating units and the Global CTO teams. The two former allow the team to ensure their efforts are technically sound and draw on the wider technical resource pool, using funding from existing access engineering budgets. With the latter, Telefónica gives full autonomy to the team on technical choices and vendor ecosystem, within an incentive and funding framework that ensures the team stays focused on the program objectives. Note that each of the three companies has involvement from senior leadership within their company or from a new business unit. Because the transformation touches so many aspects of the respective corporations, it requires strong leadership and support from the top in order to be successful.

10 Unic@ is one of the most ambitious telco transformation projects in the industry, virtualizing the company's core networks. The Unic@ program is focused on mobile core virtualization program, moving towards the other parts of the operator's network

11 For more details see <https://www.telefonica.com/en/web/press-office/-/telefonica-presents-internet-para-todos-a-collaborative-project-to-connect-the-unconnected-in-latin-america>

12 See Harvard Business Review case study: Telefonica: A Lean Elephant by Henry W. Chesbrough et al.

Figure 10: AT&T, Deutsche Telekom and Telefónica organizational set-up



6. What's new about the *new design*

Access-driven transformation replaces traditional central office aggregation function with a leaner and lower-cost design:

- CO pod is a modular pod built using many of the same components you would expect to find in a typical cloud data center, with a software stack designed for extensibility and stability in mind.
- The new design can replace a diverse range of specialist network appliances with software.
- CO pod can support network attachment, traffic aggregation, subscriber, and device management.
- The “secret sauce” of the new design is it allows protocol trimming, elimination of repeated functions and granular configuration to alter functionality to match the operations model desired by the operator.

It is widely accepted that telecoms people – from product marketers and technical managers to customers – want cloud-like agility, but what it entails is much less well-understood. Agility consists of a set of practices related to development, testing, integration, verification and deployment of software, made possible with highly structured and well-thought-out software toolchains and access to programmable and scalable infrastructure. To capture efficiencies derived from reuse, these platforms enable and encourage granular function development and their subsequent API-ization. This combination of agile practices and software reuse dramatically reduces the time it takes to get an idea into production. The result is that developers can create, launch and evolve a broad variety of differentiated services in less time than ever before.

Why the current design is challenged

Traditional telecom platforms were not designed with rapid application development or external developers in mind. These platforms are embedded systems, with proprietary vertical OSSs developed for and by vendor-developers. Their designs are typically driven through lengthy standardization processes and place a premium on service stability, operational resilience and asset life, rather than innovation. Consequently, functionality is rigid, interfaces are brittle, capacity-add is lumpy, and operational design favors operational risk reduction. All of this is far removed from the flexibility associated with the programmable, everything-as-a-service cloud paradigm. Moreover, the standards also tend to preserve the status quo. This ultimately leads to an ossified industry with a small, locked-in set of suppliers. To provide end-to-end services, OSSs are coupled through electronic and manual processes to form the “OSS/BSS mesh.” Often the mesh spans hundreds of individual applications and even spreadsheet-based processes. Tight coupling of business

and operational processes with their underlying infrastructure means that despite the complexity, the collective must evolve in tandem. In practice, this means each equipment upgrade, process or service change is conditional on its ability to reliably integrate into the mesh. This design is no accident; it is the consequence of a rational process of domain-based resource allocation in siloed organizations. But siloed responsibilities means low risk, verticalized projects with quick payback are favored, while long-term, transversal and admittedly complex projects are put off indefinitely, left to accumulate technical debt unless a profitable business case can be demonstrated. The outcome is that, as operations mature, operators become increasingly held hostage to their systems and platforms. In this environment, agility requires ingenuity in managing expectations, as well as in getting things done.

Despite the complexities, operators need their technical teams to deliver continuous improvements in customer experience, time to market, and operational efficiency. Recognizing the challenge and desiring to build a community among operators with similar goals, a group of operators funded the creation of ON.Lab to take a fresh look at the problem. Rather than working on theoretical designs, they did what software engineers do – and created a prototype software-centric solution. Building on work in open hardware done by the Open Compute Project (OCP), they positioned the central office in the image of a cloud data center. Using disaggregated routers and servers, combined with open-source software tools such as Kubernetes, Openstack, OpenNebula, ONOS and XOS, they developed a multipurpose central-office platform. Early CORD demo and field trials vetted the architecture, and ON.Lab has now transformed, in both name (now ONF) and mission, to develop production-ready systems defined by operators which are called reference designs (e.g., the broadband wireline access reference design

is SEBA). Other organizations have used the same principles to create comparable technical architectures, such as the Broadband Forum with its CloudCO program and OPNFV's Virtual Central Office project.

The new hardware architecture

The basic physical unit for the new design is a modular pod built using many of the same components you would expect to find in a typical cloud data center. It consists of a modular, high-density rack of general-purpose compute, storage, high-speed programmable switching, and special-purpose devices to enable different broadband access technologies such as FTTx, xDSL, cable or WTTN and, as of late, radio access. This design allows the pod to provide access, aggregation, edge-routing and computing services. The pod design limits the number of specialized hardware designs and maximizes the use of general-purpose OCP hardware through the disaggregation and refactoring of components that were previously found on the interior of legacy access nodes. Figure 11 shows an example implementation of a half-rack pod dimensioned to serve the fixed and mobile traffic needs of a medium-sized neighborhood of circa 30,000 households.

The new software architecture

Figure 12 illustrates high-level functional building blocks of the new design. It uses the same technology patterns that have been proven effective in public-cloud data centers. However, unlike the public cloud, the pod is a transit cloud. This means, in addition to running public cloud-like virtualized compute and storage workloads, the pod must forward traffic efficiently to


also support communications workloads. This is made possible through an architecture that explicitly enables NFV and SDN applications. The former provides infrastructure services that emulate the public cloud. Meanwhile, the latter allows developers to exploit the programmable networking fabric for data-plane traffic control and processing. This is an important step beyond traditional NFV implementations. Rather than using the compute functions for complex packet switching and forwarding, the programmable network hardware is able to offload a significant part, or even the complete data plane, of a (virtual) network function, which can reduce the amount of compute resource required. This is accomplished by using the white-box fabric switches already in place to interconnect the compute and storage. New elements are not typically needed, and software access to the white-box switches allows controller-based direction to program sophisticated behaviors into the fabric. Using the switching fabric this way creates a great deal of efficiency and cost savings compared with using general-purpose compute for switching workloads. For example, a single fabric elastically supports server interconnect, basic Ethernet switching, IPv4 and IPv6 routing, IP/MPLS switching, and even BNG and SAE gateway functions – as separate network slices – at the same time¹³.

The software elements that make up a complete implementation include both the infrastructure management and network function software. As shown in Figure 12, the former consists of open-source operations and management tooling for virtualization, configuration management, testing, monitoring, logging, analytics and security. While the architecture pattern is the same in most deployments, the specific tools used will vary and depend on the operator's preference or familiarity

Figure 11: Example hardware implementation of a converged CO pod

Example of pod edge platform

- 30,000 homes passed (30% connected)
- 21,000 mobile subscribers
- 33 4G macro sites per pod
- 5 Mbps/0.5 Mbps average peak capacity per GPON and concurrent mobile subscriber

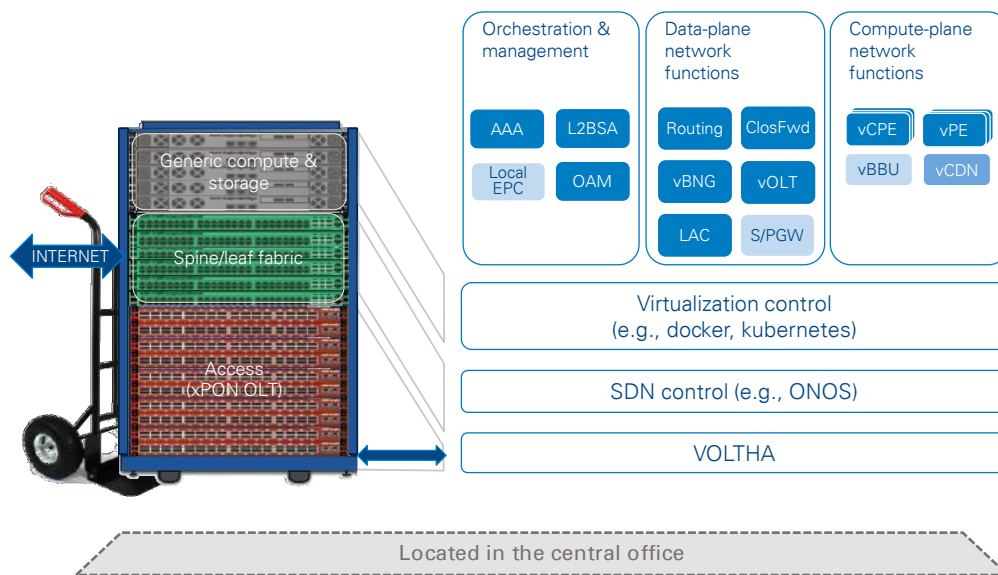


Bill of materials

- 19" 21U OCP rack
- OCP 1x32TB storage node: 8x4 TB HDD
- 4 OCP servers
 - 4x12cores 128GB RAM
 - 500 GB SDD each
- 6 OCP leaf and spine switches
 - Fabric: 32x40GE ports
 - 1.28 Tbps backplane capacity each
- 10 OCP OLT blades:
 - Fronthaul: 48x2.5 Gbps
 - Backhaul: 6x40 Gbps each

¹³ For more on the fabric, see the Trellis project: <https://www.opennetworking.org/trellis/>

Figure 12: Illustrative software implementation of a converged CO pod



from previous traditional data-center deployments. However, the choice of SDN controller has narrowed for the authors. One of the most advanced open-source options targeting a CO deployment is ONOS, in combination with SEBA and OCP design-based OLTs. As part of the CORD project, ONOS can host several other data-plane network functions, such as routing, ClosFwd, vBNG, vOLT/vOLTHA, LAC and S/PGW. The CO pod's local orchestration or management functions (AAA, L2BSA, Local EPC, OAM) and additional network functions (vCPE, vPE, vBBU, vCDN) could be hosted as either stand-alone containers/VMs or SDN controller-integrated apps. Operator deployment circumstances will drive the specific choice.

In the CO pod, VOLTHA provides hardware abstraction and vendor-neutral device models, as well as OpenFlow and CLI interfaces. The approach used by VOLTHA takes an important step beyond simply adding NETCONF interfaces and standardizing YANG models for traditional access nodes. While the latter can open management interfaces and reduce lock-in to vendor EMSs, existing standard YANG models prescribe what functions are in the hardware, and in many cases, thwart efforts to disaggregate hardware from software or unbundle complex devices into smaller, simpler elements. The CO pod is interested in software feature functions running in common compute, rather than in the access hardware. This makes changes, upgrades and vendor substitution much easier, and reduces both CAPEX and OPEX for the overall solution. To this end, VOLTHA developed a simple ethernet switch abstraction for the SDN controller, hiding differences and complexities in access PHY management as a local matter – driven by technology profiles to set the needed hidden aspects of access. VOLTHA is combined with the ONOS controller, various control applications

that embody the access feature set mentioned above, and a set of functions to provide commercial fault, configuration, accounting, performance, and security (FCAPS) management. The entire package is called SDN-Enabled Broadband Access, or SEBA. When it is combined with either open or vendor-specific hardware drivers for disaggregated OCP design OLTs or other access nodes, it allows these minimalist devices to be managed like any other access node. As a result, the access network appears as an abstracted resource with a single SDN API. Higher layer-overarching orchestration systems, such as ONAP, exploit the SEBA APIs to control the specific functions associated with disaggregated access. Routing and ClosFwd allow traffic steering and segregation (slicing) in the switching fabric. vBNG and S/PGW allow data-plane gateway functions to be offloaded and processed locally, including Internet traffic termination. Operators need to support various types of Internet access with IPoE, PPPoE and multicast capabilities. LAC implements the CO pod side for L2TP sessions that serve wholesale partner subscribers. These data-plane functions are being developed so that they, too, can be partially or fully offloaded into programmable switching hardware. Most of these data-plane functions will be driven by both an SDN component to drive control-plane functions and an orchestration or management function to drive management-plane functions. The latter can be called software defined management (SDM). For example, an AAA component uses SDN to intercept and generate typical signaling that supports subscriber authentication, authorization and accounting, which then feed the vBNG, SEBA and vEPC functional configurations. Another example is the local EPC component that manages the S/PGW. Additional functions may comprise, e.g., a vCPE, which allows a variety of services, typically embedded into customer-premise

equipment, to be brought into the network. Such platform- and component-based software architecture is designed with extensibility in mind. Each function can be extended or complemented with additional functionality that can be deployed in the same platform.

The overall stack is designed to provide stability and availability through carefully crafted, software-based designs that reduce risk through modular software components, which can be subject to granular analysis to isolate technical issues. Moreover, clearly delimited geographic scope provides a safe place to learn and experiment with the infrastructure management and network function software. This provides operators with an environment where teams can learn new ways of working from agile, DevOps and site reliability engineering.

Workloads supported

The CO pod replaces a diverse range of specialist network appliances with software, as shown in Figure 13. The pod does not duplicate the functionality of this equipment verbatim; rather, it uses the fact that individual functionality can be stitched together to rewrite the overall operational model. Specialist appliances, in many cases, have wider functional capabilities and may have more energy- and footprint-efficient designs. However, the flexibility benefits of software platforms outweigh the static benefits of higher-performance, but rigid, hardware-based platforms.

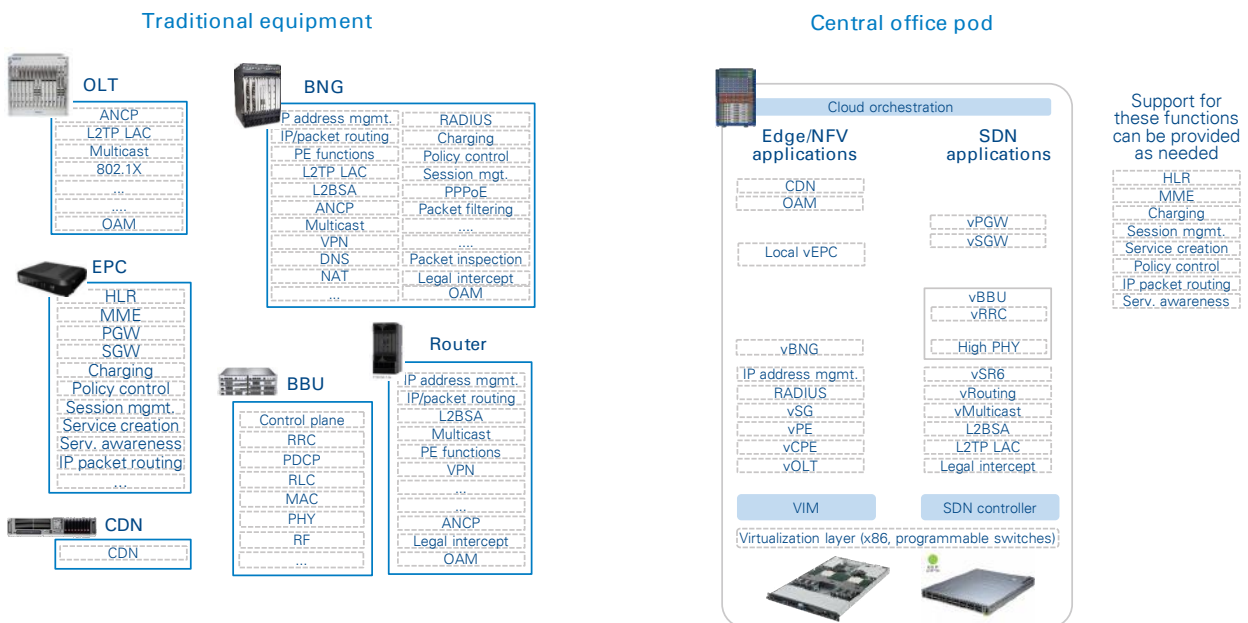
An appropriately equipped pod can support multiple workload scenarios. Within the fixed network, the CO pod can support network attachment, traffic aggregation, subscriber, and device management. However, it can also become a services platform – taking on edge functions such as vCPE, vCDN and LAN-based

security services such as parental control, firewalling and virus detection. In mobile networks, the CO pod can support edge functions such as vBBU for split RAN deployments, localized S/P-GW, and other applications traditionally provided at packet core locations. Bringing mobile and fixed network functions together into a “hybrid pod” allows true convergence, which enables access-agnostic traffic aggregation and management. All sorts of services can be delivered equally to any sort of access that the pod can support: residential or business; wireless or wireline. The CO pod can also be used for much more than hosting operator edge-network functions, addressing cloud value pools. With the right security and isolation between internal and third-party workloads, spare capacity can be made available to external partners. Taking advantage of the locational proximity of the pod to provide low-latency edge-computational services allows latency-sensitive workloads to be deployed. Examples include augmented reality, localized data-intensive workload processing, and eventually ultra-reliable, low-latency services such as those described in 5G use cases.

The secret sauce

Moving from traditional to disaggregated architectures makes three important process and service changes possible. First, traditional telecoms equipment is interconnected with standardized interfaces and protocols that often require long chains with bits and pieces of value added along the way. Virtualized programmable access networks can be used to reduce these chains by collecting all the value from different physical locations with the centralized SDN controller applications or CO pod orchestration and management component (see inset: “Trimming protocols”). Since the controller applications are open, they can be customized to directly access the central databases that typically host the

Figure 13: Functional comparison of traditional equipment and converged CO pod



policies and service attributes for each customer, device, and circuit. As a result, an SDN-controlled CO pod has less internal protocol complexity and is much easier to troubleshoot and modify. Second, repeated functions can be eliminated. Most traditional network elements contain aggregation, routing, and additional, often unique functions such as policy enforcement or a unique physical layer. In the CO pod, the fabric provides aggregation and routing for all local functions, so those local functions can be simpler and more specialized. Separate policy enforcement and physical layer functions are integrated with the fabric – and thereby gain common aggregation and routing capabilities. Finally, because the functions in programmable access networks are open and more granular than those in traditional network elements, it is possible to eliminate unwanted functions, as well as alter functionality to match the operator's desired operations model.

Much of the OPEX incurred by operators is involved in mapping the operational paradigm and capabilities the equipment providers have envisioned to those employed in their own networks. Open source can be altered to become plug-and-play and match the operations model where it is used (see inset: "Open-source & operations code"). Each aspect of the service that can be automated reduces operations costs, including expensive truck rolls. Moreover, each hop saved saves equipment, power and associated labor costs. At scale, individual tasks convert into millions of dollars. The same story of de-duplication along the service chain yields similar energy savings.

Emerging NFVi and VNF software ecosystems are not yet on par with data-center equipment, which means the industry must navigate its way towards one or more credible solutions. As a key component of virtualized access networks, programmable switching hardware allows offloading of packet traffic to devices that are tailored for providing high throughput at low cost and energy consumption. Programmable ASICs have made significant progress and allow integration of programmable hardware, such as the network elements in the switching fabric as VNF offloading functions. This goes hand in hand with the separation of the control and user planes and the functional decomposition addressed at ONF and 3GPP, the latter for mobile networks. However, it is unlikely that one size will fit all due to differing production focus and skills. AT&T, Deutsche Telekom and Telefónica, individually and collectively (as part of ONF), are working on making the CO pod production ready. While the scope and priorities of their activities vary, they are all trying to figure out the right commercial and open-source software toolset to give operators adequate technology control to accelerate change, as shown in Figure 14. For this reason, the operators that participate at ONF have made substantial changes in the governance and goals of the organization. The new approach focuses on deliverables that operators are willing

Trimming protocols

Two examples of simplifying protocol chains are found in multicast distribution and attachment and authentication (AAA).

Presently, multicast distribution is provided box by box, with a daisy chain of signaling that installs multicast state in each device. Each multicast-enabled box must support the same standardized application and variant of protocol used among boxes. Moreover, the protocol must be made robust against lost signaling by periodically re-syncing the state of the entire multicast tree. As a single SDN domain, the CO pod only needs to support multicast signaling at its boundaries: to the customer and the upstream network – usually with IGMP and PIM, respectively. The SDN controller hosts a multicast application that instructs boxes at these edges to send inbound signaling messages to itself. The multicast application knows the distribution state of the entire pod, and when new requests require changes, it instructs the proper boxes to adjust their multicast filters. No protocols are required in the interior of the pod, and no applications are needed in every element.

For AAA, there is a similar chain of protocols and device-specific applications from the CPE to an AAA database, which supports authentication, attachment, accounting, and applying common service policies. Like with multicast, the protocol from the CPE and that from the AAA database is directed to the SDN AAA app. The app also collects other information, such as which access port is requesting attachment and which instance of policy enforcement is associated with that port. Finally, using SDN can create service topology on the fly, eliminating the need to pre-provision service circuits. The SDN AAA app queries the database or its local cache directly for policy and AAA data.

to deploy. It brings the work farther from the demo and trial software maturity level – and closer to an operationally complete and deployment-ready level.

Finding open-source software ecosystems is not difficult; the real task is to find a reasonably small number of communities to which operators actively contribute. These contributions cannot just be requirements or standard specifications but need to be in the *coin of the realm* of open-source communities: code.


Despite the challenges, the CO pod design provides an elegant foundation to catapult the industry production model into the future. However, more importantly, it unlocks a world of possibilities to transform the operator.

Open-source & operations code

Open-source functions, when available, are a boon for operators. Unlike traditional network elements, in which the functions are an amalgam to serve many customers, according to a model envisioned by the equipment maker, open-source functions can be modified to work according to the operational model preferred by the operator. An example in which changing such a function yields simpler operations can be found in the vOLT SDN application for broadband access. When a new customer device (ONT) is discovered on a PON by vOLT, it is attached to the PON and then compared to previously seen devices. If it was seen before, nothing more need be done, and if it is brand new, a side process is

launched to bind the device to a customer account when that customer authenticates. Compare this with the prevailing process today, which requires a technician or end user to input/provision the serial number or registration code of a device into the OLT system before it will allow that device to work on the network. It should be clear that by changing the process flow in this software, an OPEX-heavy process can be replaced with a lightweight, automated one. The benefit is similar to that which plug-and-play technology brought to attaching peripherals to computers. Moreover, a carrier can make this change to open source, where previously it might have had to propose new standards and gain agreement from many other companies to support such a change.

Figure 14: CO pod hardware and software component alternatives



Category	Component	Alternative 1	Alternative 2	Alternative 3	Alternative 4
Hardware	Compute	2 x QTC QuantaGrid D51B-1U	Various	3 x HP 2 cores with 128GB RAM 2x 2Tb HDD each 1 NIC 40GE each	2 x Open CloudServer: 4 socket 12 cores, 128GB RAM, 2 x 500GB SSD NIC 2x40GE each
	Storage	none	Virtual using SSD in Compute	Virtual using disks in Compute	JBOD: 8 Disks of 4TB => 32 TB
	Networking	Spine 2 x OCP Accton 6712 switch Leaf 4 x OCP Accton 6712 switch	Various OCP	Spine 2 x OCP switch Leaf 4 x OCP switch	Spine 2 x CloudEngine 7800 DC switches Leaf 4 x OCP Accton 6712
	Optical access	Edgecore, Adtran, Celestica, Ciena	OCP, based on AT&T spec	OCP based (existing and DT contributed specs)	Celestica
Software	OS	Ubuntu LTS 14.04	Ubuntu/ONL	Ubuntu LTS/ONIE/ONL	Ubuntu
	VIM	OpenStack	Akraiho (Kubernetes)	Docker/Kubernetes	Open Nebula
	SDN controller	ONOS	ONOS	Homegrown/ONOS	ONOS
	Service management	XOS	ONAP, OSAM	Homegrown Orchestration & NEMO	OneFlow
	Monitoring		ONAP	ELK	Iris Kibana
	vRouter	Quagga	DANOS	rtbrick	vErC (IPv6)
	vSwitch	trellis	None - All-silicon user plane	rtbrick	CLOSfwd
	RADIUS	802.1X, DHCP	802.1X – LDAP	For AAA (PPP subscriber facing)	none
	vOLT	VOLTHA	vOLT/VOLTHA	VOLTHA	VOLTHA/Celestica
	vSG		Various (v)BNG	-	vPdC1 (IPv6)
	BBU		-	-	Altiostar
	EPC functions		-	-	Quortus, Athonet

7. How the central office pod changes everything

The CO pod considerably widens choices for operator production platforms development:

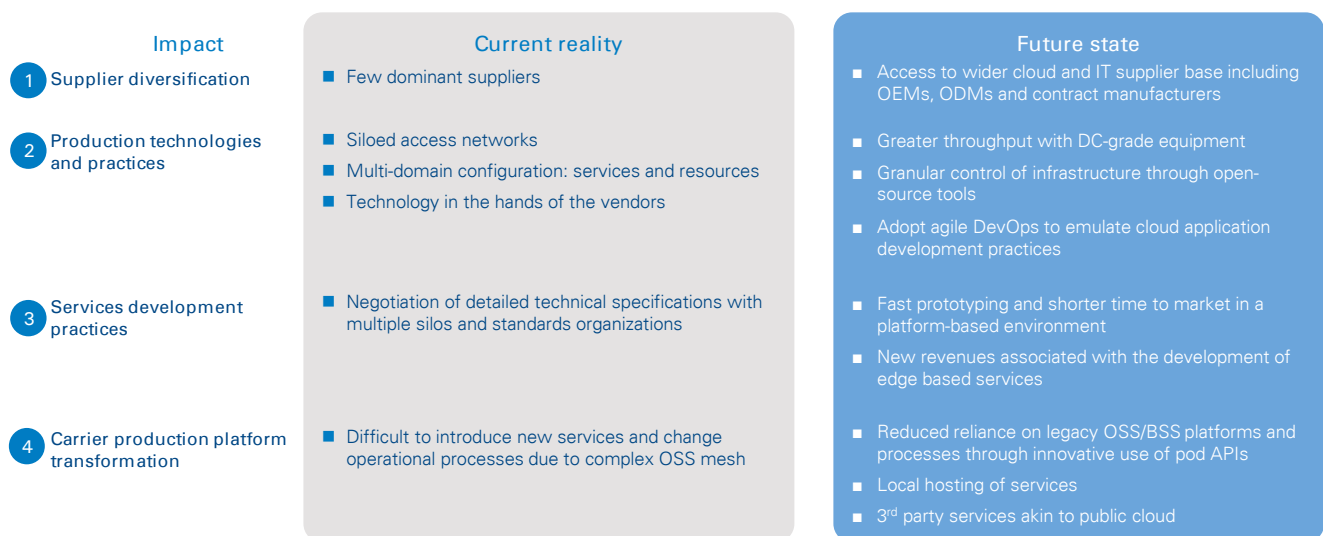
- Use of commodity data center-grade hardware provides an alternative to concentration in the telco supplier ecosystem.
- An operating model based on highly automated, open-source software tooling allows operators to emulate cloud service providers' innovation and economics.
- Multipurpose infrastructure empowers product marketers and service engineering managers to rewrite rules for innovation.
- The locality factor can be used to discontinue legacy processes.

The CO pod-based design affords operators the luxury to broaden the telco equipment ecosystem. On the face of it, the pod substitutes the traditional central-office aggregation function with a leaner, lower-cost and non-blocking design. However, it is much, much more, as illustrated in Figure 15, and opens up many new strategic development options.

The **first major impact** is supplier diversification and monolithic equipment disaggregation. Powerful procurement functions, as well as industry competitive dynamics, have resulted in concentration of the telco-supplier ecosystem in fewer network

elements that represent large, bundled, single-source solutions to typical network siloed verticals. Traffic growing faster than revenue means finding lower-cost alternatives is a priority. Use of commodity data center-grade (e.g., OCP Accepted¹⁴) hardware provides such an option. This equipment allows the industry to benefit from throughput-oriented equipment designs that are free of obligatory operating systems, control applications, EMSs, and vendor maintenance contracts. Equivalent white boxes are available from multiple suppliers and can be integrated with the control and operations software that an operator desires. More importantly, use of data-center

Figure 15: Impact of access network virtualization



¹⁴ To become OCP Accepted, a white box must be open, available to anyone, and provide schematics, board layouts, and even firmware and FPGA/CPLD code. An OCP Accepted spec can be taken by any manufacturer and turned directly into that OCP product

designs brings the economies of scale of the cloud service provider supply base into the telco network. In addition, the white-box approach enables better assessment of the real value of each hardware and software component, which limits the ability to create opaque bundles. The first impact replays an ecosystem change that has already happened in data centers. In the past, computers, operating systems, and even applications were sold as vertical bundles. Today no one would buy a server that only supported one operating system or forced them to buy the applications from the same vendor. This is how we see the future telecom ecosystem as well.

The **second major impact** is to allow operators to emulate cloud service providers by deploying similar production platforms and operations paradigms. At the core of a web-scale cloud platform is an operating model based on highly automated, open-source software tooling, which provides granular control of infrastructure and tools to track everything. The tools are contributed line by line by independent developers and cloud behemoths. Data center service providers have been striving toward “lights-out” facilities, and end users that make use of cloud resources cannot be disappointed because they cannot properly manage and operate their workloads. Thus, data-center automation is a well-refined, mature set of operations and processes that take even the smallest details into account. In many data centers, entire racks are placed and retired with no human having ever touched the equipment within those racks. Cloud-scale hardware, combined with scale-hardened open-source tools, puts cloud-like operational economics within reach of the telco industry. As shown above, reframing services into cloud workloads and architecture patterns allows dramatic simplification of fixed- and mobile-access network components. The complexity removed from these components is de-duplicated, stripped to its bare essentials, and pushed into the orchestration and management of those components, which, in turn, are pushed into the highly automated systems developed to manage the cloud. Aligning production architectures with the cloud brings the industry technically on par with cloud service providers. It also allows operators to take back technological control of their production platforms, along with end-to-end responsibility for the platform, including when “things go wrong.”

The **third major impact** is the ability to change how services are developed. Availability of multipurpose infrastructure enables the central office to become a focal point for new and varied workloads. Using agile software development practices, combined with infrastructure and network services available in the pod and exposed to third parties as APIs, allows product ideas to be put into practice quickly and built upon by many. The ability to *just do it!* in a platform environment empowers

product marketers and service engineering managers to rewrite rules for innovation. New-product development can be in-sourced, out-sourced, even crowd-sourced. Consider this in comparison to typical new-product development in telcos today. First, there is negotiation of detailed technical specifications in minute detail with multiple organizational silos from IT, network, operations, procurement and finance. Then the results drive new standards and features, and sometimes even new network elements developed over 12–24 months. Then an RFP is issued, and eventually a vendor is selected that can deliver a real service in another six to nine months. Often by this time, the product requirements have changed significantly, or an OTT has developed and deployed that product and taken first-mover advantage across many operators. In the new ecosystem, product stakeholders can jump into the driver’s seat, prototype new services themselves, and run those services on existing infrastructure in their networks – and potentially in other networks around the world that offer the same infrastructure-as-a-service (IaaS). With the right platform and engineering capabilities to support them, development efforts are structured as discrete micro-services enabling software reuse with open APIs. As the catalog of (micro) services increases, more software reuse becomes possible, and the time to create a new working demo is progressively reduced, from months to weeks to days. In this future, the entire product development process will turn on its head. Roles will need to be redefined around building of features in successive sprints to deliver working prototypes, minimum viable products, and incremental product enhancements.

What was just described could be done with any cloud, so what is different about the CO pod? The pod is a delimited infrastructure resource only loosely coupled to the workloads it supports, which means it can be managed, provisioned, upgraded and orchestrated in complete isolation from the rest of the network. Moreover, if vCPE is employed, the pod can host applications that exist in the customer’s LAN, as opposed to some distant place in the public cloud.

As a result, the pod can support novel services for various customer devices, as well as manage the customer’s LAN traffic (e.g., improve wi-fi). This allows operators to offer more intimate functionality within their customers’ networks, improve network experience, and expand service to include device management. All of this can create new value for customers (e.g., parental control, IoT device management, nearby storage and back-up, software-defined enterprise-branch connectivity, and in-home traffic prioritization). While by no means exclusive, access to this data can unlock numerous value pools in the security and smart-home spaces. But that is not all. Under the right conditions, local application hosting can be extended to third parties, in a

way similar to the public cloud, using an edge-infrastructure platform service model to reposition the operator in order to capture additional revenues in the business of public edge-cloud provider.

The outcome is an opportunity to rethink product innovation in a delimited location. This allows each operator to experiment and test new product ideas and software tools at their own pace, prior to widespread deployment. However, taking advantage of the platform requires profound changes in how services are created and deployed, and will require rebuilding of core technology skills and the vendor ecosystem. (This topic is discussed in greater detail later in the paper.) While this requires change, it is not as daunting as it may seem, because these new skill sets are abundant in today's marketplace; they overlap with the skills needed for modern IT and web-scale systems. In many ways, the change should make it easier to find skilled staff than it is today to find people with traditional telecom skills.

The **fourth major impact** is the ability to use the central office as an internal production platform to accelerate operator transformation. Clever use of the locality factor can be used to discontinue legacy processes and thin-down BSSs/OSSs and legacy service delivery platforms. Exploiting the fact that the customer and the serving pod are on the net, local hosting of service delivery platforms with direct interaction becomes attractive. Such thinking creates new degrees of freedom and provides new ways to deal with the technology and operations hurdles found in some operators' back offices.

The combined effect of the above is that the CO pod changes the economics and risk profile for **service and production innovation**. It does this through providing a safe place to experiment and learn quickly, even with real customers. Once the design is tested, validated and hardened to carrier grade, powerful software distribution tooling allows it to be rolled out almost instantaneously across all pods. This also helps to reduce the minimum workable scale at which a service can be profitably offered. However, for the CO pod to become a reality, the industry must face, head-on, the operationalization challenges that will be described in the next section. Lastly, because services live on common infrastructure, there is no sunk capital investment to try new services. Operators can be less risk adverse, try new services with less certainty about their popularity, and simply repurpose the infrastructure if the service is a flop.

8. Operationalizing the central office pod

To make the new design a reality, operators must address four operationalization challenges:

- Replacing end-to-end integration vendors with a white-box vendor proposition that is limited to design, manufacture, shipping, reverse logistics and repair.
- Finding where to home the CO pod, recognizing that power and cooling needs may be higher than is accommodated in a typical central office.
- How to approach pod insertion into the production landscape: a new greenfield design or bootstrap into the existing landscape.
- Dealing with the complexities of pod deployment and service cutover without the support of the incumbent vendor community.


Despite the simplicity of the CO pod hardware and software stacks, making cloud technologies telco-production ready is a big step and out of the typical operator’s comfort zone. It will require willingness to let go of industry handrails because there is not yet a credible roadmap for putting cloud technologies to work in an operator network.

We discuss four key operationalization challenges, summarized in Figure 16. Many aspects will need to be worked through, including how to source the pod, where to deploy it, technical integration of the platform, and who will deploy the platform into production. These decisions will shape how the technology can be used to create business value.

CO pod hardware and software sourcing

The first CO pod operationalization challenge is *sourcing the hardware and software*. This should be trivial, but it may not be, due to prior practices at operators. For new production environments, the industry modus operandi is to buy services from soup-to-nuts solution providers. However, going forward, there is a desire to avoid lock-in, select best-in-class solution components, and optimize performance. Doing this requires combining products from numerous white-box or grey-box hardware vendors, as well as commercial and open-source infrastructure management tools, in a coherent platform engineered for carrier-grade reliability (see inset: “Of white and

Figure 16: CO pod operationalization



gray boxes”). It’s likely that both niche vendors and system integration vendors will be needed, as well as traditional telco vendors. Additionally, because of the smaller scope of their contribution, many niche vendors are unable or unwilling to offer the breadth of services and assume the implementation risk typically sourced from traditional vendors. While traditional vendors often can provide such integrated solutions, their approach to doing so typically involves maximizing use of their proprietary products and services.

Of white and gray boxes

This paper provides extensive discussion of white box, grey box, and black box. For the purposes of this paper, a white box is a non-differentiated, generic implementation of hardware. It typically is fully decoupled from software, and the two are sourced more or less independently. So white boxes from different suppliers are essentially interchangeable, and they are not tightly coupled to any sort of software or optics. White boxes are typically specified openly, such as those described at the Open Compute Project. A gray box is like a white box but provided from an OEM with a brand attached. These typically have an OS provided by the same OEM but can often also support other OSs. Black boxes are those provided as vertically integrated platforms of hardware, software, and support – for which only one (typically OEM) supplier provides everything.

Because white boxes are decoupled from software, and there can be many best-of-breed contributors to an overall solution, more suppliers must get involved in the new ecosystem. These include sub-components, integrators of sub-components, and overall system integrators. In a world where system and security patches are ongoing, each of these entities has an ongoing commitment to the overall solution.

The base white-box vendor proposition includes design, manufacture, shipping, reverse logistics and repair. Software integrators’ business models cannot assume liability for software patching and security issues. It requires someone to take on the responsibility of coordinating diverse specialties, from rack equipment and white-box vendors, to hardware and software integrators, to deployment and testing vendors, all while taking care of materials flows and logistics. As if the burden was not enough, someone must also take care of managing the life cycle of each element of the platform while it remains in production.

While it is plausible that one or more end-to-end integration vendors could emerge, the most realistic option today is for operators themselves to create and rebuild skills and assume end-to-end responsibility of the new stack. If openness is desired, the operators need to have the technical and business capability to change their solutions from one white-box or open-source solution to another. Operators must onboard white-box vendors and figure out the support model that works mutually, in addition to taking care of materials flows and logistics. An operator’s decision to commit to state-of-the-art, open-source tools can be forestalled by making use of commercial tools or commercially supported open-source tools. Doing this can also help the operator gain support to climb the learning curve in order to use the tools directly. Operators can contract third parties to pre-integrate the pod, including hardware assembly, racking, wiring, and software integration test and deployment. However, the last of these could also be internalized. Reskilling operator staff to perform these tasks using modern e-learning tools and step-by-step guides may create goodwill and lower the total cost, if it fits within ordinary operational activities.

CO pod placement

The second operationalization challenge is to determine *where to physically install* the CO pod. This might not seem obvious, but the CO pod might not always be in the traditional central office space, which has implications for the design of physical architecture of the pod and for the set of local codes and rules that apply. The CO pod is based on an OCP-compliant, standard 19-inch rack that has to operate from -42.5V to 72V DC power. While IEC/ETSI central-office racks have the same width, they are typically much smaller and shallower. In addition to the physical differences, central offices might not have the weight or environmental performance of a typical IT data center or elevated/raised floors for cooling. The central office is a temperature-controlled location that supports the environmental class ETSI 3.1, and thus places an upper limit on power density per rack at 8 kW. Installing more equipment in an already-overcrowded central-office facility may require additional capital spend on power systems and environmental conditioning. How this plays out depends on whether the operator is an incumbent that owns the central offices or a challenger that does not. Lastly, even within the CO, there are multiple fit-for-purpose spaces that may be selected, including traditional telecoms space, co-location space, and space that may have been upgraded to support typical data-center guidelines. The TIA is working to describe these options, as well as advance and clarify requirements for CO transformations, and OCP has recently defined a new OpenEdge ecosystem that is designed specifically to be deployable at cell sites and in traditional CO space.

Challenger operators that struggle to secure appropriate space in existing central offices may find it cheaper and more agile to place all or part of their infrastructure into other operators' co-location space or Internet peering facilities. In addition to the data center-grade space to deploy unmodified OCP racks, such a move would put challengers in prime position for peering with Tier 1 web-service players. In contrast, incumbents will likely deploy in central offices at marginal cost, if space is available. However, if significant capital investments are required to upgrade facilities, the shift to a professionally managed data-center facility might make more sense. The final choice depends substantially on bilateral negotiations and regulatory regime, as well as the depth of the peering facilities available in each geography. The outcome has important implications for the design of the CO pod: low capacity may encourage splitting the design into access components and disparate central components managed by regional cloud, whereas high capacity might favor a converged rack design in single locations.

CO pod technical and operational integration

The third challenge is figuring out the right approach *for insertion of the pod into the network*. To date, introducing a new network element has meant plugging the device into the operator OSS mesh. But the CO pod is very different; its frontline positional advantage, combined with its ability to be used as a platform, creates additional deployment options.

The most obvious deployment option is abstracting the CO pod platform to appear as a traditional hardware appliance, hence limiting the operational impact of introducing the platform into existing provisioning, monitoring and assurance processes. In effect, this uses modern tools to re-create legacy control functions to simplify integration. However, an alternate approach can use the fact that the CO pod is the "last" set of equipment towards the customer. The pod can be used to selectively or completely bypass existing or legacy OSS/BSS/omnichannel applications and processes, as well as service delivery platforms, without disrupting the OSS/BSS management chain. Doing this can create new degrees of freedom to innovate and deal with technology and operations problems.

CO pod roll-out and service cut-over

The fourth challenge is operationalizing *deployment and service cut-over*. Deploying the CO pod requires figuring out how to get it into production without causing significant customer service disruption. Deployment services for traditional telecom solutions are highly structured and typically provided as turnkey services that bundle equipment, project management, installation, commissioning, test and cut-over. However, this service does not exist for the CO pod, and there is no standardized playbook

on how the entire process should be managed. Nonetheless, it represents a considerable endeavor requiring new skills to minimize cost and service disruption.

One approach is to train and develop existing telecoms suppliers to support the transition to the new environment. An alternate approach is to internalize the process by reskilling field staff to perform integration as well as deployment, making the activities compatible with daily operational activities and moving their workflow toward DevOps. While this might slow deployment, it could create a stronger foundation for platform evolution from initial to target architecture.

Decisions and restrictions around how to source, integrate and deploy determine overall program scope and complexity. Our view is that there is no right answer today on where to deploy and how to insert the platform into production. However, arguably, the hardest challenge of all is not the above; it is the status quo bias. To influence the direction and pace of technical evolution, operators must learn the importance from their cloud peers of taking control of their own technology fates. In practice, this means building world-class cloud engineering capabilities and a supplier ecosystem to match, as we describe in the next section.

9. Building service and platform engineering capabilities

Success requires operators to acquire new skills in engineering and supply chain and rethink investment planning:

- Supply chain sourcing functions accustomed to predefined end-to-end specifications must evolve.
- Operators must learn the art of using low-cost, general-purpose hardware and virtualized software to (re-)produce carrier-grade products and services.
- Engineering capabilities will require a concerted effort to attract and cultivate the right skills.
- Operational use of open-source requires cultivating (and investing in) community projects with other operators and competitors.

In a world of disaggregated hardware and software, operators must acquire new skills and change attitudes towards technology made outside the operator ecosystem. Building these skills will require focus and operating outside of the organizational comfort zone. The difference between success and failure will include the ability to develop a minimum critical mass of in-house engineering capabilities, as well as to expand the supplier ecosystem.

Changing engineering skill set

Operators must learn the art of using low-cost, general-purpose hardware and virtualized software to (re-)produce carrier-grade products and services (see inset: “Operating virtualized infrastructure”). That does not mean operators must build everything themselves. The required functionality can be developed by combining services to achieve the desired features, including services offered by the platforms, third-party service APIs, and small and discrete, validated pieces of in-house software. In cloud architectures, performance and service availability depends on infrastructure software engineering, rather than on high-end boxes customized to deliver “five nines.” These aspects will often be unrelated to the supplier of a discrete tool or function. Hence, the responsibility for overall system design cannot be transferred to a third party; rather, it lies squarely in the hands of the operator. Ultimately, this means operators must become self-reliant, taking the responsibility of defining, engineering, implementing, commissioning, and testing their own implemented solutions. To execute this vision, operators must develop skills that include managing orchestration, cloud networking and large-scale distributed

Operating virtualized infrastructure

In virtualized environments, applications (tenants) must consider security, processing scheduling priority, and operations management. In early operational models at AT&T, this led to an increased number of “admin” staff. This was because early operations models did not alter the way the staff organized their work. Operations tools developed in cloud ecosystems can be used to simplify telecom deployment and management. The new virtualized environments have built-in automation that is not available in vendor-siloed solutions from a few years ago. For example, to update software, you make a policy change in the orchestrator, and that sets automation tools in action to accomplish the update. If they are not adequate, rollbacks are just as easy because they are achieved through the same automation. This method is much easier than managing scripts and libraries and cycling through all the instances to ensure they are the same across the board. It also results in less people time, less processing time and more accurate and predictable implementations. Once AT&T teams were reorganized around the tooling available for the virtualized environment, the number of “admins” dramatically reduced. These benefits were also extensible to managing hardware. Typical OCP specification hardware has APIs that make it much easier to load loosely coupled operating systems and other firmware images. This supports life-cycling hardware and software independently, increasing agility and resilience.

The central office pod for technologists

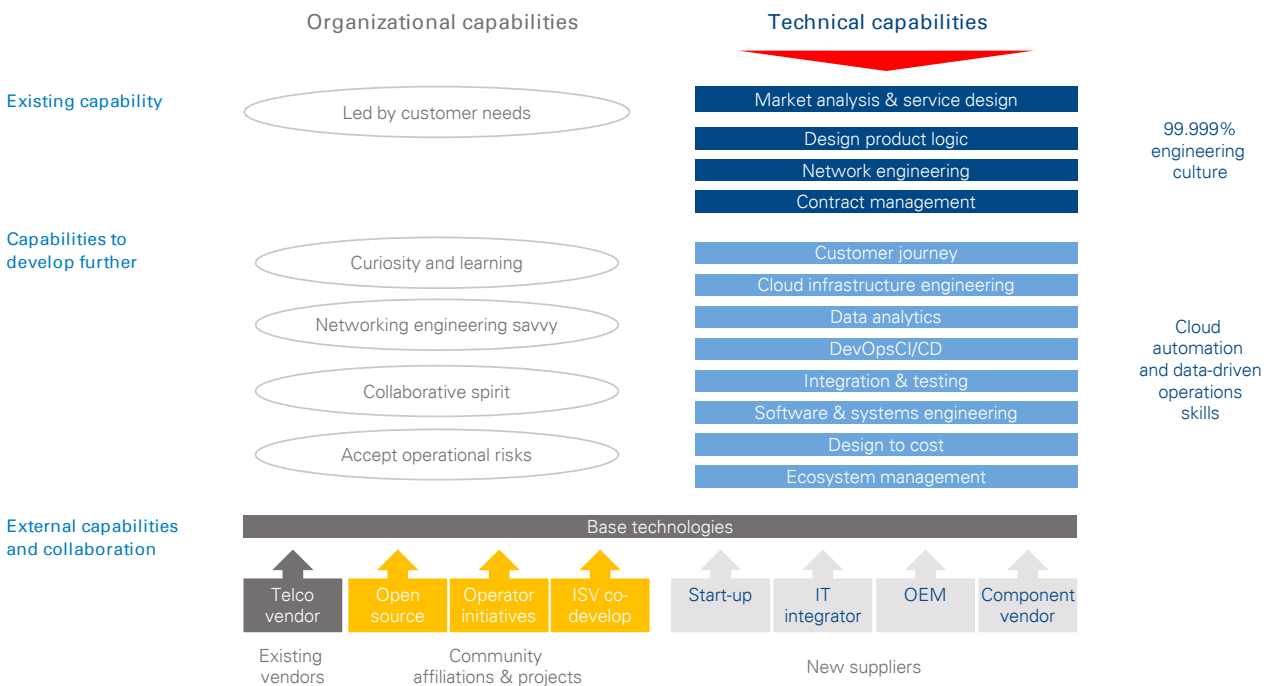
systems engineering. Smaller operators might rely on larger peers to lead the effort of developing standard designs and use system integrators to deploy solutions based on these designs. Regardless of their size, we think operators will benefit from engaging in open technical communities to develop and mature these capabilities, share learnings, and discover best practices.

Because of the potential disruption to their business models, operators cannot rely on whole-hearted support from traditional telco vendors to adopt cloud architectures; they must expand the supplier ecosystem. To de-risk and accelerate the expansion, they must find alternative ways to onboard cloud technologies. This means looking beyond existing suppliers, employing new models of collaboration, and fostering community-based mutual support. Cloud hardware suppliers have a very different relationship with their customers than telco suppliers have with theirs. Cloud hardware providers typically supply system components, with the responsibility for overall integration lying with the customer, which allows them to relentlessly focus on component excellence. The focus on components means they engage more intimately with component vendors. And in the more general case, their limited scope of activities means that they are unlikely to be able to offer the services the operators have become accustomed to when working with traditional vendors, and also that they have lower cost structure for the services they provide. One example of such a service is maintaining inventory so operators do not need quarterly forecasts and month-long lead times for manufacture. Another is maintaining staff to provide product support, training, and education. While this continues to be the case for smaller suppliers, we are seeing larger ODM suppliers create subsidiaries that help close this gap.

Cloud tooling is typically open source, but access to the code is not a warranty that it is operationally robust. Successful operational use of open source requires much more than downloading code from GitHub; it is about community. It entails willingness to invest time and resources for the benefit of the community, as well as authorship. If operators are to benefit from these technologies, each must find an approach compatible with its scale and technical competence. Seeking commercial support from a tier-one community member might appear to be simplest, but in a cloud environment with many moving parts, it is also the hardest to operationalize – it is on par with traditional vendor support. The flipside is a DIY approach that requires proactively investing in multiple open-source communities and ecosystems. It also involves attracting active project contributors, along with their knowledge and expertise. This inevitably changes the split of labor and responsibility between in-house teams and third-party vendors.

This environment, shown again in Figure 17, will look very different from that of the past. It will combine traditional vendors and ways of working with new vendors that engage with very different contractual relationships and ways of working. It will be a heterogeneous ecosystem with responsibilities thinly spread among many, rather than the few. The shift to this approach will require many new skills in multiple functions: architecture and engineering, operations, procurement and finance organizations. While this may seem unattainable at first, we believe every operator can do this, as AT&T, Deutsche Telekom and Telefónica have already succeeded in this endeavor. The most important task is to onboard all stakeholder functions from day one and set an appropriate pace that allows everyone to remain engaged.

Figure 17: Capability and ecosystem map



Widening engineering capabilities will require a concerted effort to attract and cultivate the right skills. The wider virtualization movement provides important lessons. Attracting and retaining talent must go beyond creating isolated workspaces for multidisciplinary collaboration. It requires offering compelling, long-term career propositions, which means being able to articulate how changes in the telecoms industry will create new and exciting opportunities for professional development. The same is true for talented individuals that are willing to work outside their own comfort zones, as well as the corporate comfort ones. With the right framing, training, and incentives, they can substantially contribute to success. Activating these individuals allows teams to tap into corporate insight and personal relationships to get things done. Talent is willing to move and invest in themselves if they feel they are part of a winning team of other talented people. However, this involves progressive changes to organization, pay scales and staff development practices. This is an art, more than a science, which requires a fundamental change in attitudes towards staffing.

The change has a major impact on operations. Lines between external and internal responsibilities become fuzzy, but the pressure to keep uptime remains unchanged. The cloud DevOps movement addresses this issue head-on. It provides guidance on how to create focus around site reliability, based on the software engineering and behavioral aspects of the team. It provides recipes on how to ensure that service continuity does not become a victim of service innovation. Today's highly siloed operations from access, transport, service platforms and security give way to a multidisciplinary operations team that includes software developers. Enabled by automation and analytics tools, these teams can rapidly deploy, roll back and problem solve, ensuring that end-to-end service reliability management is just as agile as writing code. These practices fundamentally change the trouble-ticketing and escalation approach to issue management in favor of team-based responsibility. However, this can create challenges for staff that are unwilling or unable to make the transition.

Impact on the supply chain model

Supply chain sourcing functions accustomed to predefined, end-to-end specifications must change. The new normal is to build what you need, combining one or more pieces of open-source and commercial software on general-purpose hardware to deliver the required functionality, and buying components from a range of suppliers rather than large, turnkey systems.

The consequences of this change are multifold:

- Sourcing processes must adapt to identifying and engaging with small innovative suppliers.

- Multi-geography and end-to-end, turnkey contracts will give way to dozens of small vendors that do not have the scale to service operators akin to existing suppliers.
- Procurement must enhance due diligence on small suppliers to ensure new vendors do not create undue operations risk.
- For cash-constrained disrupters, the operator may pay a greater portion of the development up front or seek additional business opportunities for the disrupter.
- In cases in which there is concern about ability to scale, an operator might phase roll-outs or use an integrator or other established company to serve as intermediary.
- An operator can promote capabilities to other service providers via press releases, roadshows, or other means if they feel that the supplier has too much dependence on a single customer.
- Procurement must place a premium on acquiring or in-sourcing know-how from smaller vendors, in addition to their components.
- To ensure better alignment of objectives and development of a healthy supplier base, procurement will need to innovate contractual models to share benefits, as well as risks.

Working in collaboration with engineering, supply chain sourcing must understand which technologies are key and take on the challenge of expanding the supplier ecosystem. This means close collaboration on engineering objectives and roadmaps. But beware; small, innovative vendors need special consideration and processes to limit the disruptive effect of the traditional vendor engagement process. Moreover, because they lack scale, small vendors may never be able to support the traditional vendor process, and this can change the procurement model permanently.

Finally, finance functions will no longer focus efforts on approving "siloed" business cases. Instead, they will use detailed asset utilization data to approve incremental hardware and software investments. In a production platform built on shared hardware- and software-based functions, business cases and budgeting will give way to competitiveness and demand.

While getting started requires a credible story, it can be done. And in the next section we show how, by describing the AT&T, Deutsche Telekom and Telefónica processes to develop the right engineering and supplier capabilities.

Learning new tricks is hard. Knowing how to start is often harder. This section provides insight from the technical teams at AT&T, Deutsche Telekom and Telefónica on how they have implemented the process of creating breakthrough solutions.

10. The technologist's guide to getting started

There is no such thing as the "best approach": program design and monitoring must be tightly linked to corporate KPIs:




- AT&T's initiative is structured around access and the overarching Domain 2.0 program.
- Deutsche Telekom AG is following an engineering approach; its primary goal is to redesign access networks to drive "step-change" reduction in life-cycle costs.
- Telefonica's program revolves around a virtual company to create focus around generating new revenues.

Their journeys are works in progress, but nonetheless, they provide important insight into program framing and structuring, as well as guidance on several other topics that are important for execution. This section is intended as a guide for technologists on how to structure similar programs and teams, as well as monitor progress.

As shown in Figure 18, AT&T and Deutsche Telekom have collocated the teams within their operating unit engineering organizations. Meanwhile, the Telefónica initiative is structured as a virtual company with funding from Telefónica's Innovation function. The former allows the teams to execute holistic approaches to integrating colleagues from engineering, design, planning, operations, the field technician team and lifecycle management to ensure their efforts are technically sound. The latter allows Telefónica to provide full autonomy to the team on technical choices and vendor ecosystem, within an incentive and funding framework that ensures the team stays focused on

the program objectives. As we will show, key to their success is creating a supportive environment that allows learning, experimentation and collaboration with third parties to address clearly articulated problems. Another success factor has been fostering a culture of continuously questioning and challenging the solution from a cost, feasibility and maintainability perspective.

Figure 18: Description of access transformation programs

			
Project design	"Not a trial! Real transformation of technology, skills, business process, and operations"	"Retain deep in-house understanding of all solution components, to allow rapid prototyping"	"Simplify, open and <i>just do it!</i> "
Description	<ul style="list-style-type: none"> ■ Structured as a corporate-wide project with engagement from Wireless and Wireline engineering and Network Infrastructure and Operations teams ■ Significant work done in communities ■ Open source and open specs first, merchant silicon and disaggregation ■ Design the system as plug-n-play and self-healing, and to be simple for operations and OSS 	<ul style="list-style-type: none"> ■ Uses the "design to cost" (DtC) principles for product & service development, i.e., actively managing the cost-to-value ratio from day one ■ Planned DevOps principle ■ Design based on commodity/open hardware and software where possible ■ Utilize partners' expertise for hardening and production-grade development, while leveraging open-source and community efforts ("collaborate & win with partners") 	<ul style="list-style-type: none"> ■ "Internal start-up" structure with a complete multidisciplinary team (marketing, strategy, development, operations, etc.) ■ OnLife Networks' primary focus is identifying and developing new revenue sources from access network virtualization and edge computing ■ Doctrine: radical simplification, no legacy, open source whenever possible, always open interfaces, do it yourself before buying



AT&T took the lead in launching the first telco transformation program focused on pivoting the network to software in 2013. Led by its Technology and Operations organization, with direct oversight by the officers, the Domain 2.0 (“D2.0”) program was conceived. D2.0 covers mobile core and enterprise network domains, focusing on customer self-service and speeding up service time to market (see inset: “From ECOMP to ONAP”).

Domain 2.0 Virtual Access (“D2VA”) is a component of D2.0 focused on disaggregating and virtualizing (and eventually converging) access networks. D2VA is built on a new approach of composing virtual network functions in a cloud-native way, which allows them to control merchant silicon. The adaptation of SDN principles to merchant silicon, especially silicon that supports wireline and wireless access technologies such as 5G and PON, is a critical enabler for bringing virtualization and automation to the edge of the network. Rather than being developed internally within D2.0, this idea was pursued and matured as an open community endeavor through ON.Lab (later Open Network Foundation – ONF). It led to various proofs of concept, demonstrations, and field trials of CORD, and then additional flavors of CORD to address specific vertical access segments. As CORD has matured, AT&T has followed the technology, as well as creating additional, complementary open-source projects to facilitate the maturation of VNFs toward cloud-native architectures, containerization, micro-services architecture, and instrumentation. Some of the resulting projects include *DANOS* as a white-box routing platform, *Akraino* as an edge-cloud deployment automation platform, and *Airship* as an edge-cloud life-cycle management platform.

The D2VA team is organized into three major workstreams: Infrastructure, Wireline and Mobility. The overarching priority for all workstreams is to mature toward an industrialized platform. As Mobility matures, the vision is for the work to simplify into two workstreams: infrastructure and converged access – with a strong search for new workstreams (business opportunities such as those described earlier in this paper) at the various edges of the network.

The first work stream defines the D2VA technical approach to infrastructure. The D2VA work stream is aligned with other D2.0 infrastructure programs but is working on a somewhat different problem. Its goal is to produce a small, hardened edge-pod design that can support multiple access technologies. The approach is very different from AT&T Network Cloud’s small distributed data centers or typical AT&T super-scale data centers, which have raised-floor, high-density power racks and highly controlled environmental conditions. The D2VA design

From ECOMP to ONAP

As an early adopter of NFV, AT&T learned quickly that there were no existing systems or defined architectures for managing and maintaining VNFs that also eliminated the EMS functions that often locked suppliers into the IT stack. AT&T believed new cloud infrastructure and VNFs were not optimally managed by classic OSS/EMS stacks or MANO, but rather, needed to be managed and packaged in more cloud-native ways. This led AT&T to develop a global service automation system called ECOMP (Enhanced Control, Orchestration, Management and Policy). ECOMP supports virtualization in VM constructs and is built on infrastructure managed by OpenStack. Rather than retaining classic OSSs and EMSs, it allows development of services in high-level compositional tools, and then maps the service requirements to capabilities from both legacy and VNF functional entities. AT&T also grew appreciative of the benefits of open source and open specs – especially in the power of community endeavors, so ECOMP was eventually brought to the Linux Foundation and merged with Open-O, and the resultant system is now called Open Network Automation Platform (ONAP). Today, over 70 percent of mobile subscribers worldwide are supported by carriers who have adopted ONAP.

must operate with concrete floors, low-density power racks and modest cooling: the environment of typical central offices. The AT&T design relies on variants of the Open Compute Project designs, and OCP suppliers have supplied equipment that can comply with safety and environmental standards of central offices.

The second stream, Wireline, is where the bulk of D2VA investment has been focused as of March 2019. The solution had matured enough to move from PoCs and trials towards field deployment. Hence, this program is transitioning toward being part of *business as usual*, managed by the *Converged Access and Device Technology* business unit under a common DevOps model for the entire business. The initial work was done at the Atlanta Foundry, which is the focal point for much prototyping work and provides the right environment to collaborate with third parties. The Foundry team is composed of innovators, software developers, and telco access experts, working under one roof. This allows D2VA to benefit from the vast experience and know-how of AT&T and has helped the entire organization flip the switch to developing and deploying access, so that the new method has become business as usual. This has helped transition and restructure this part of the operator into a DevOps model, and follows similar transitions to those that have already occurred for enterprise and infrastructure capabilities. The D2VA Wireline team is actively

involved in ONF. Rather than building software in a silo, they participate in community efforts and with other operators at ONF. The community has adopted a largely similar architecture to meet the software needs of wireline access and calls the project SDN-Enabled Broadband Access (SEBA). Several operators have agreed to make SEBA the core of their similar deployment plans and hope to enjoy benefits from community and ecosystem development. At the time of this writing, AT&T is engaged with the other authors of this paper, as well as other collaborators at the ONF, to develop and deploy this system.

The third work stream is focused on mobility. This work stream seeks to develop a consensus architecture among operators for disaggregating the mobile RAN through collaboration in the O-RAN Alliance community. Like the Wireline work, the wireless team seeks to progress the R&D as a community and operates through the O-RAN, ONF and other communities. Activities for initial deployments are underway.

All workstreams are focused on pivoting the network to software, which is key to unlocking the business imperatives declared in the Domain 2.0 vision white paper: Open, Simplify, Scale. These imperatives are pursued in several software-focused activities:

- **Disaggregation:** Separates the software logic from the hardware directed to perform it. It's often called CUPS and is the underlying original intent of SDN.
- **Orchestration:** Uses common software tooling over and over for each workload and infrastructure, rather than deploying a stovepipe EMS, NMS, and OSS for each box and service.
- **Automation:** Gaining access to open-service logic through disaggregation allows automating the service logic itself (to become plug and play), as well as automating the management and operations of that service using powerful emerging technologies such as artificial intelligence and machine learning.
- **Self-service:** Software-driven customer portals can allow customers to help themselves and gain superior experience with a supplier, but this term also means exposing underlying capabilities *as a service* so developments that come later can reuse that capability with little friction or tight coupling. It is often called *everything-as-a-service*.
- **Collaboration:** No one – not AT&T, not even the huge legacy telco suppliers – can solve it all, write it all, or know it all. The telco transformation is a very large undertaking and depends on many collaborating to get it done. Nowhere is collaboration more supported than in open-source and open-spec communities, and in those communities, *code* is the coin of the realm.

AT&T has realized that while it may have the resources to create functional software prototypes, it lacks scale and capabilities to maintain them as ongoing, proprietary, in-house systems. Following that ethos, every member of the D2VA program is encouraged to engage, participate and develop community through contributions, staffing events and bilateral collaboration or “co-creation” with other single entities.



Deutsche Telekom is taking a re-engineering approach to its access virtualization program. A4.0's primary goal is to redesign access networks to drive “step-change” reduction in life-cycle costs. Its mandate covers fixed and mobile access networks, as well as edge-computing infrastructure, but initial focus is on FTTH, FTTN and mobile backhaul.

The program uses “design to cost” (DtC) principles for product development and manufacturing. DtC has been practiced at Deutsche Telekom since 2003, inspired by earlier work at a major German car manufacturer. DtC is a systematic approach to controlling the cost-to-value ratio of product development and manufacturing. A central theme in DtC is that costs end up being designed “into the product” at an early stage and are difficult to remove later due to over- or misinterpreted requirements. Hence, they ultimately produce ill-conceived solution designs. Key DtC steps include baseline cost determination, discovery of cost elements that drive total system cost, cost-trend analysis, cost avoidance option identification, and action and implementation plan. Applying DtC, the A4.0 team's technical objective was to produce, at the lowest-possible production cost, a design based on commodity hardware and software that would support any existing service transparently. Here, transparency means the design will work and be managed just like the network elements it replaces. Focusing on transparent access brings additional OSS/IT complexity but has allowed the A4.0 team to operate with little involvement from the product or commercial functions. A4.0-driven DtC should not be seen as intent by Deutsche Telekom to develop or manufacture telecoms equipment; rather, it is used to drive community-based specification discussions with current and new vendors to drive win-win exchanges.

In a hack lab, the A4.0 team dissects a range of different types of telecoms equipment – OLTs, BNG, edge routers, RAN and mobile gateways – to identify sources of cost, cost drivers, and “bills of materials”. Moreover, they have also been exploring alternate approaches to streamline the embedded software landscape. This has allowed the team to develop a “map of components and services,” as well as identify areas of redefinition and simplification, and how pieces should be

sourced (in-house, community or vendors) without losing carrier grade. The A4.0 team has yet to define the details of the target *modus operandi*, but recognizes that commercial software, as well as integration work, will continue to play a material role. The current thinking is that software that drives features visible to customers or defines the internal business model should remain in-house, and all else can be sourced from the community or proprietary vendors. However, Deutsche Telekom must retain skills to take third-party code, look into it, find issues (add and remove parts of it), and debug and test it. However, the burden of turning it into production-grade systems should fall on organizations with strong software integration skills, commercial or otherwise. Where these software components are not specific to operational processes or value-add, A4.0 envisages the use of supported open source; however, for service-specific components, various shades of proprietary systems are also envisaged.

Execution requires building a talent pool. The initial team included about 30 people in a cross-functional team from Deutsche Telekom Technik Germany and the Deutsche Telekom Group, and specialized in access engineering, IP networking, OSS, security, planning, operations and service assurance. These were complemented with software engineers from third parties. However, as the program moved from ideation to development, missing skills were supplemented through partnering arrangements. The current thinking is that, rather than building a full set of capabilities in-house, A4.0 must rely on partners with deep software engineering skills. In order to co-develop software with partners, the A4.0 team hired additional software engineers with skills in architecture of large-scale systems, data science, engineering and container networking. The team also relies on skills from external sources, as well as ONF capabilities.

Funding and progress monitoring is facilitated by the fact that the program sits within a single business unit. Program progress is monitored by systematically tracking the quality of the vendor ecosystem willing to support A4.0 roll-out, and by continuously updating the DtC cost-estimate model.

Telefonica

Telefónica launched its network transformation program, Unic@, in 2013. Unic@ is one of the most ambitious telco transformation projects in the industry, virtualizing the company's core networks. From ideation to design, proofs of concept, definition and launch, Unic@ is now deployed in four countries. The dramatic effort of the company has created a telco cloud model underpinned by SDN and NFV. Recognizing early on that orchestration was a key success factor of telco cloud creation,

the company open sourced its open MANO program through the ETSI open-source MANO (OSM) effort. The Unic@ program started as a mobile core virtualization program, moving towards the other parts of the operator's network.

Onlife™ Networks was born outside of Unic@ in May 2016, under the auspices of the innovation team. Though Onlife™ has many possible development angles, its primary focus is on identifying and developing new revenue sources from converged access networks and edge computing. After successful completion of an initial stage gate, prototyping was approved in July 2016. Thereafter, Onlife™ went through the process of techno-economic verification until field trials were authorized in September 2017 and first commercial clients were connected in June 2018.

Onlife™ sits within the Networks Innovation group, alongside other disruptive product innovation units. It operates as an internal start-up with its own executive team, as well as services and development teams. The Onlife™ project is organized around four disciplines: Infrastructure, Access, Platform, and Services. The Infrastructure team is responsible for all aspects related to access, switching and compute infrastructure, as well as virtualization software or the VIM stack. The Platform team defines and develops APIs for internal teams and third parties, and is responsible for catalog definition, service orchestration and billing. The Services team is responsible for business development, as well as supporting services definition, prioritization and testing.

Over two-thirds of the team are developers with skills in access and cloud infrastructure engineering. The remaining staff are technical managers and product managers. The team operates in three locations: R&D centers in Madrid and Valladolid, as well as a test lab in Madrid; each facility has its own testbed. Minimum viable product focus has meant Onlife™ team has worked on producing product prototypes, rather than validating multiple vendors. The technical team has been responsible for 100 percent of the integration efforts. They have selectively adapted open-source tools for their needs (ONOS, SEBA) in conjunction with open-source VIM vendor OpenNebula and vBBU vendor Altiostar. Hardware has been sourced from Celestica, Tibit, and Edgecore, and integrated by Flex. As the product moves closer to industrialization, rather than building the capability in-house, the Telefonica team expects to seek external support. All Telefonica Innovation projects have designated sponsors. The role of the sponsor is to support the project through the ideation, prototyping, industrialization and productization stages. Onlife™ sponsors include the UNIC@ team and the Spanish business. The former provides technical guidance, whereas the latter has supplied the production environment to test the solution. As part of the latter, Onlife™ is in close collaboration

with the Spanish business network engineering teams to trial use of the platform as a substitute for triple-play, as well as edge use cases focused on gaming and enhanced video entertainment. Future development plans include scaling the solution up to over 500 residential clients and 10 enterprises, deploying multiple use cases on the platform in a production environment with revenue clients, and validating the platform and willingness to pay through A/B testing.

Progress monitoring and accountability are based on progression through the “Lean Elephant™” stage stage-gate process. At the current stage of development, the relevant metrics are focused on solution productization and getting the platform into production with real customers and scale. These include metrics for the number of residential and business users

connected, services/use cases, and deployments across the Telefonica footprint.

Table 1 summarizes the insight from the work done at AT&T, Deutsche Telekom and Telefónica for newcomers. Their experience shows the variety of issues and solutions to problems that are typically taken for granted when using operator ecosystem equipment. There is no one-size-fits-all approach to these challenges, but nonetheless, the proposals highlight possible approaches. One of the more difficult issues will be deciding which commercial and open-source software solutions to bet on. At this stage, it is our view that there is no right or wrong; however, that should not be an excuse to do nothing.

Table 1: Insight and lessons learned

	Problems and issues	Possible solutions
Architecture/ engineering	<ul style="list-style-type: none"> Key technology components are not mature, or still being developed and rapidly changed. Integrate cloud and networking: Most network processes are hardware based, whereas new paradigm is software based. CI/CD for infrastructure and applications are essential for operations, but usually not part of standard vendor solutions. Automation of deployment is critical (infrastructure, platforms and applications). OS and software-hardening know-how is limited. Need to implement sophisticated identity management for service operators’ roles and policies. Engineering faced with transformation – technical and professional resistance against designing for MVP with cyclic upgrades versus traditional waterfall methods. 	<ul style="list-style-type: none"> Avoid early technology lock-in and test different alternatives. Implement DevOps methodology and be ready to develop in-house pieces of the solutions. Develop and deploy virtualization (cloud) and networking (SDN) components as a whole solution, avoiding making them separate entities to be integrated in the future. Deploy CI/CD environments and automatize everything. Decouple POD lifecycle from OSS/BSS. Implement firewall, traffic characterization, logs storage, encrypted communication, VPN solutions. Develop guidelines for security configuration, best practices, patching policy, protected software repository. Need for clarity on roles and policy associated with each role implemented directly in the system.
Operations	<ul style="list-style-type: none"> Double “load”: Technical and professional resistance against continuing support of legacy platforms and processes while ramping up skills for the future mode of operation. Large gap between existing operations model versus cloud: stability/few changes versus components constantly changing due to CICD. Security teams need to understand “cloud-native” processes. 	<ul style="list-style-type: none"> Solve real problems and show it is better than legacy. 80% of focus should be on people, processes and behavior. Align fragmented responsibilities: Plan/build/run to DevSecOps → start with small scope, grow iteratively. Embed life-cycle management concept for all software, including cloud environment.
Testing	<ul style="list-style-type: none"> Current testing and certification/verification processes are not well-suited for software-based networking and cloud production. Testing is not a continuous process. E2E testing requires integration with legacy systems & access. 	<ul style="list-style-type: none"> CI/CD environments required for infrastructure, platforms and applications testing, E2E test automation implementation is a must. Need to adopt agile methodology and work iteratively towards target picture. Avoid integration with legacy as much and as long as possible.
People	<ul style="list-style-type: none"> May lack training for the new approach. Difficulties to “sell” using open-source and in-house development versus vendor solutions. Fear of “build, test and learn” cycle in network teams. Culture and mind-set of silo-thinking from line managers and lack of risk-taking attitude from top management (used to offloading risk to vendors). 	<ul style="list-style-type: none"> Agile training, promote transformation, establish incentives. Dissemination of (small) “success stories”, seek support from exec management. Select companies supporting open source commercially. Strong community. Mistakes are not penalized, but opportunities for development.
Partners	<ul style="list-style-type: none"> Partners are facing similar challenges, few true “cloud-ified” products in portfolio. Supply chain transformation to “software services, including integration” ongoing – not fully ready yet. Cooperation models are nascent. Challenge to work with small partners/start-ups. 	<ul style="list-style-type: none"> Challenge status quo to enable partner cooperation, use open-source and actively develop an ecosystem. Trial with new partners for network software and SI. Actively participate in telco-specific open communities and seek collaborations. Work through larger partners (sub-contracting).
Legal and finance	<ul style="list-style-type: none"> Customers and their data belong to the operating unit and handling them raises significant concerns. The operating unit is expected to keep traceability and interception of traffic as requested by authorities. Complex and costly solutions are involved. 	<ul style="list-style-type: none"> Create a GDPR profile and negotiate a parent company/operating unit legal agreement – clearly establish each one’s responsibilities – enable a trust relationship this way. For limited field trials, exceptions may apply. For a larger base of friendly users, the trial should be integrated with existing systems on a per-service basis (e.g., IP traffic, voice).

Closing remarks

Key messages:

- Regenerating technology, capabilities and supplier ecosystems based on web-scale technologies is a strong opportunity.
- The CO pod provides a new way to deliver results quickly: distributed architecture provides a safe place that allows learning on the job.
- No operator can or ought to go it alone; the industry must “do it together” to drive lasting change.
- In the new future, operators must become more self-reliant and take greater responsibility and leadership in the value chain, while pragmatically challenging the status quo bias.

Like it or not – and for good reasons – operators, their suppliers and standards bodies are facing major change. Collectively, the industry must regenerate its technology, capabilities and ecosystems to deliver competitive solutions based on web-scale design approaches and methodologies. Moreover, we believe this change is a strong opportunity for operators, as well as new entrants to the telco ecosystem.

In this paper, we have made the case for access disaggregation and softwarization via the CO pod. The collective view of the authors is that it provides a new toolset for faster and broader operator transformation. Enabled by its front-line location, the CO pod provides operators with a low-cost, safe place to virtualize existing and test new service concepts, using well-understood cloud methodologies. Moreover, distributed deployment means more resilient services and a smaller span of outages. It also enables learning by doing, without the risks of multi-domain transformation programs.

However, unbundling the industry-standard access technology stack means operators must figure out how to reassemble and execute, and who to rely on for help. It is our view that no operator can do it alone; working collectively, operators need to look beyond traditional ecosystems to attract new talent and vendors that thrive on speed and innovation. However, sufficient talent and selection of vendors may not yet exist, so each operator individually and the industry in general must cultivate them. Operators must also enlist entrepreneurs and entrepreneurial capital to join their cause. Moreover, they must have the self-confidence needed to provide leadership to existing ecosystems and create new ones when needed. Ultimately, this will encourage new forms of collaboration, renewal and business-model innovation, which will allow the

industry to expand and disentangle from its current situation without creating new traps along the way. The long game includes creating spaces where innovation can and does happen. In this future, international standard bodies (e.g., ITU, 3GPP) and industry associations will continue to exist, but reference designs and exemplar platforms born from collaboration in open-source communities are going to play a key role in setting de facto standards for real-world implementations.

With the right culture of talent, pricing models and partners, success is all but guaranteed. But success will require putting skin in the game, which requires commitment to create a supportive environment that will allow experimentation and focus on meaningful business challenges. This means operators must become more self-reliant and take greater responsibility in the value chain to define, develop, integrate, test and commission in-house solutions to drive differentiation and pragmatically challenge the status quo.

Remember, Who Dares Wins: If you are willing to dare, take the leap and join the community – you are unlikely to win if you don't!

Notes

Notes

Notes

Glossary

3GPP	3rd Generation Partnership Project	LAN	Local area network
AAA	Authentication, authorization and accounting	NETCONF	A protocol defined by the IETF to “install, manipulate, and delete the configuration of network devices
API	Application programming interface	OAM	Operations, administration and management
BBU	Base band unit	OEC	Open edge computing. See www.openedgecomputing.org
BNG	Broadband network gateway	OCP	Open Compute Project
BSS	Business supporting systems	OLT	Optical line termination
CDN	Content delivery network	ONAP	Open Networking Automation Platform. See www.onap.org
CLI	Command line	ONF	Open Networking Foundation
ClosFwd	ONOS open-source application for segregating traffic	ONOS	Open Network Operating System. See: www.opennetworking.org/onos/
CO	Central office	OpenFlow	A communications protocol that gives access to the forwarding plane of a network switch or router over the network
CO pod	Central office pod	OSS	Operating support systems
CORD	Central Office Re-architected as a Datacenter. See www.opennetworking.org/cord/	PE	Provider edge
CPE	Customer premises equipment	PHY	Abbreviation for the physical layer. A PHY connects a link layer device (often called MAC as an acronym for medium access control) to a physical medium such as an optical fiber or copper cable
CUPS	Control and user-plane separation	PPPoE	Point-to-Point Protocol over Ethernet
DC	Data center	PON	Passive optical network
DTBen	Deutsche Telekom Basic European Network	PGW	Packet gateway
EMS	Element management system	SEBA	SDN Enabled Broadband Access. See www.opennetworking.org/seba/
EPC	Evolved packet core	SGW	Servicing gateway
FCAPS	Fault, Configuration, Accounting, Performance, and Security	SDN	Software-defined network
FTTx	Fiber-to-the-home, building, cabinet, etc.	TIP	Telecom infrastructure project. See telecominfraproject.com/
ITU	International Telecommunications Union	VM	Virtual machine
IPoE	IP over Ethernet	VOLTHA	Virtual-Optical Line Termination Hardware Abstraction www.opennetworking.org/voltha/
L2BSA	Layer 2-Bitstream Access	XOS	ONF’s VIM and service manager
L2TP	Layer 2 Tunneling Protocol	YANG	(Yet Another Next Generation) is a data modeling language for the definition of data sent over the NETCONF network configuration protocol
LAC	L2TP Access Concentrator		



Who Dares Wins!

How access transformation can fast-track evolution of operator production platforms

Arthur D. Little

Arthur D. Little has been at the forefront of innovation since 1886. We are an acknowledged thought leader in linking strategy, innovation and transformation in technology-intensive and converging industries. We navigate our clients through changing business ecosystems to uncover new growth opportunities. We enable our clients to build innovation capabilities and transform their organizations.

Our consultants have strong practical industry experience combined with excellent knowledge of key trends and dynamics. ADL is present in the most important business centers around the world. We are proud to serve most of the Fortune 1000 companies, in addition to other leading firms and public sector organizations.

For further information please visit www.adlittle.com or www.adl.com.

AT&T

We help family, friends and neighbors connect in meaningful ways every day. From the first phone call 140+ years ago to mobile video streaming, we innovate to improve lives. We have the nation's fastest wireless network.** And according to America's biggest test, we have the nation's best wireless network.*** We're building FirstNet just for first responders and creating next-generation mobile 5G. With a range of TV and video products, we deliver entertainment people love to talk about. Our smart, highly secure solutions serve nearly 3 million global businesses – nearly all of the Fortune 1000. And worldwide, our spirit of service drives employees to give back to their communities.

The AT&T Technology Organization is a world-renowned leader in working with developers and start-ups to open our network to innovation. AT&T is one of the largest R&D investors in the telecoms business. In addition, AT&T is a founding member of the ETSI Industry Specification Group (ISG) on NFV as well as ONF, and supports numerous industry open-source projects from Open Compute Project (OCP) and Open Network Foundation (ONF), and Linux Foundation projects including ONAP, O-RAN, AKRAINO and DANOS.

AT&T Communications is part of AT&T Inc. ([NYSE:T](https://www.nyse.com/quote/NYSE:T)). Learn more at att.com/CommunicationsNews.

Deutsche Telekom

Deutsche Telekom is one of the world's leading integrated telecommunications companies, with some 178 million mobile customers, 28 million fixed-network lines, and 20 million broadband lines, across more than 50 countries. With a staff of some 216,000 employees throughout the world, Deutsche Telekom generated revenue of 75,7 billion Euros in the 2018 financial year.

DTAG has been instrumental in encouraging the development of standards and software communities to drive softwarization of the industry. DTAG has been a launch partner of ETSI Industry Specification Group (ISG) on NFV, OpenCompute, ONF, and TIP. DTAG is also active in encouraging collaborative innovation through multiple vehicles, such as its partner-level membership at the "new" ONF, as well as collaborative research with academia. Hub:raum is a start-up incubator. Deutsche Telekom Capital Partners has invested in dozens of innovative technology companies which are intended to complement the business activities of DTAG. DTAG's engineering divisions specialize in bringing actual innovation to the live network. More recently, to address a gap in the market, DTAG has created MobileEdgeX, a collaborative effort with other operators to create a technology platform for mobile edge computing.

For further information, please visit www.telekom.com/.

Telefónica

Telefónica is one of the largest telecommunications companies in the world by market capitalization and number of customers, with a comprehensive offering of world-class fixed, mobile and broadband connectivity services. The company has a significant presence in 16 countries, with strong presence in Spain, Europe and Latin America. In 2018 Telefónica had consolidated revenues of 47.8 billion Euros, consisting of 270 million mobile phone subscribers, 13 million Internet and data lines, and 8 million pay-TV customers. Telefónica is a 100 percent listed company, with more than 1.5 million direct shareholders. Its shares are traded on the Spanish Stock Market and on those in London, New York, Lima, and Buenos Aires.

Telefónica ("TEF") is a prominent member of the first wave of operators that recognize the transformative potential of using cloud technologies. It is understood that using general-purpose hardware with virtualized network functions could transform the capabilities and revenue-generating potential of the network. Moreover, it could operate the network more cost-efficiently and reduce complexity and time to market for new service development and launch. To support this vision, TEF has invested considerable management and R&D resources in its telco cloud, Unic@, and the associated NFV reference lab, as well as development of an ETSI-compliant multi-domain orchestrator platform open-source MANO, available as an open-source tool for the community. In addition, Telefónica is an active voice in industry forums and other open-source communities, such as ETSI NFV, ONF and TIP.

For further information, please visit www.telefonica.com/.

** Based on analysis by Ookla® of Speedtest Intelligence® data average download speeds for Q2 2019

*** GWS OneScore, September 2019

AT&T products and services are provided or offered by subsidiaries and affiliates of AT&T Inc. under the AT&T brand and not by AT&T Inc. Additional information about AT&T products and services is available at about.att.com. Follow our news on Twitter at @ATT, on Facebook at [facebook.com/att](https://www.facebook.com/att) and on YouTube at [youtube.com/att](https://www.youtube.com/att)

©2019 AT&T Intellectual Property. All rights reserved. AT&T, the Globe logo and other marks are trademarks and service marks of AT&T Intellectual Property and/or AT&T affiliated companies. All other marks contained herein are the property of their respective owners.