Telco Cloud
Executive Briefing Service

JUNIPER NETWORKS

STL PARTNERS

Executive Briefing

TELCO CLOUD: WHY IT HASN'T DELIVERED, AND WHAT MUST CHANGE FOR 5G

Telco cloud made big promises for the transformation of telecoms. It is a fundamental enabler for 5G and the exciting opportunities ahead. Why hasn't it delivered yet – and what needs to change?

Matt Pooley, Practice Lead | matt.pooley@stlpartners.com | February 2020

Executive summary

Telco cloud: big promises, undelivered

Telco cloud, which we define as deployment of virtualised and programmable telecoms infrastructure (NFV, SDN, AI, automation, distributed computing, and more) to support adoption of the innovative business practices used by successful cloud computing companies, was supposed to deliver big:

Figure 1: The telco cloud promise: Major buckets of value to be unlocked

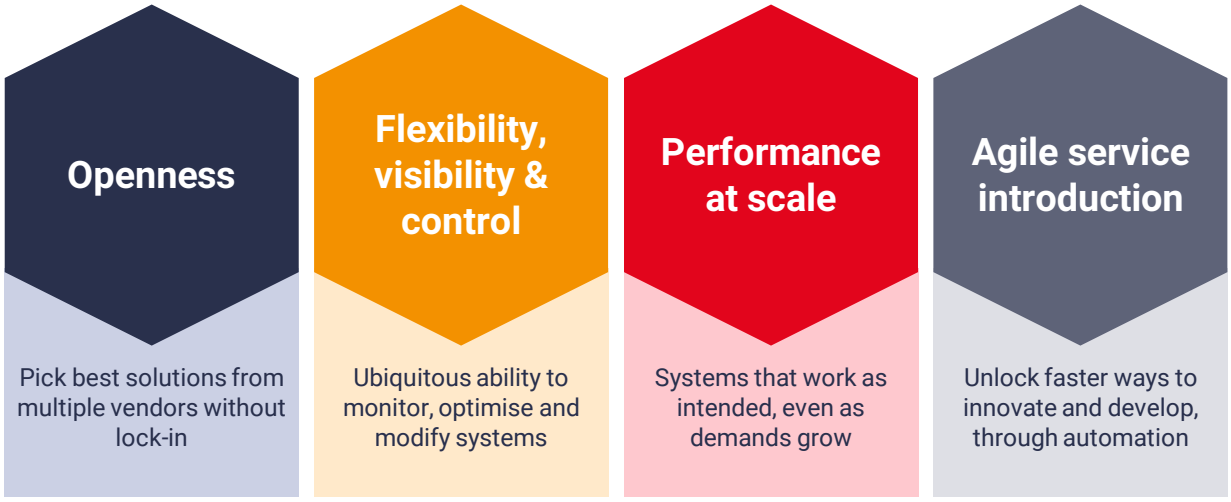

Source: STL Partners

Telco cloud technology is starting to be adopted by telecoms operators across the world, primarily driven by the move to virtualise mobile core networks in response to data traffic growth, and in preparation for roll-out of 5G networks. To date, most of it is based on Network Functions Virtualisation (NFV).

Despite a surge in deployments, operators we work with are increasingly frustrated in the results. Solutions to date are far from open and vendor-neutral. The ability to monitor, optimise and modify systems is far from ubiquitous. Performance is acceptable, but nothing to write home about, and not yet proven at mass scale. Examples of truly innovative services built on telco cloud platforms are few and far between.

The problem: flawed approaches to deployment

Our analysis of discussions with more than twenty operators worldwide has shown us that most telco cloud deployments to date fall into two categories, driven by the operator's starting point in making the decision to proceed.

Three-quarters of operators have taken a **function-first** approach to building telco cloud, driven by an identified urgent need to virtualise their Evolved Packet Core, which handles data on 4G mobile

networks. This will help to deal with data traffic growth, as well as keep the door open for deployment of virtualised 5G cores in future.

Most of these operators have deployed what we term a *vertically-integrated NFV stack*, in which a single vendor (usually a network equipment provider) provides all or most of the network functions, infrastructure and MANO capabilities, bundled in with the vendor's own computing equipment.

This "one-stop shop", which comes with a comprehensive service-level agreement (SLA), can be rolled-out quickly. However, vertically-integrated stacks tend to be closed ecosystems. As a result, operators coming at telco cloud from a function-first standpoint have realised very little of the openness, flexibility, visibility and control over their networks that they were promised. Overall, we hear that a vertically-integrated approach is successful in the short term, but in the longer term, the cons outweigh the pros.

Some operators have taken a **platform-first** approach. They aim to build a single cloud platform that can be used to support any and all telco applications, now and in the future. Components from different vendors can run side-by-side and swapped in and out at the operator's wish. Through doing so, operators can foster greater competition between vendors, driving performance improvements and ultimately delivering a better end-user experience.

The problem lies in the lack of a practical pre-existing blueprint for a telco cloud platform of this kind. Operators are forced to take a do-it-yourself approach to making it happen. Building a multi-vendor ecosystem, with components that play nicely with one another, is far from straightforward. The operators which have succeeded, most of whom are multinational groups, have been able to throw significant resource at defining, building and maintaining their telco cloud platform. Most operators simply do not have the financial or human resource to do this successfully.

The solution: change, collaboration and integration

From all of our discussions with operators, however, we find that a **multi-vendor telco cloud platform *is* the preferred goal**. Those that have come at NFV from a function-first perspective have seen the drawbacks of working with a small pool of vendors that cannot offer a truly "cloud native" environment. Those who have attempted to build platforms continue to believe that this is the best approach to take to deliver on the telco cloud promise. Consensus is that if the telecoms industry can overcome the barriers to taking this approach, we will be well on our way to unlocking the full potential of telco cloud.

Problems are as follows:

- **The need for internal transformation:** Telco cloud represents a wholesale transformation of a telecoms operator's infrastructure and operations. As such, it requires a wholesale rethink of how the organisation works. Whole organisations must "swallow the magic pill" and commit to making change happen, through leadership, acquiring new skillsets, and quantifying what success should look like.

- **The need to foster collaboration and integration:** For a multi-vendor, multi-component telco cloud platform to work, we need to ensure that all the elements – new and legacy – can co-exist harmoniously. None of this is achievable without significant cross-industry collaboration

- **The need to balance standards and blueprints:** ETSI's NFV Industry Specification Group has been the driving force of such efforts for telco cloud. But operators tell us that while standards are important, they lag behind their realities and needs. As such efforts are emerging to take ETSI's work a step further and define reusable, actionable reference models which take the best aspects of the standards, but leave the operator and vendor partners free to innovate within it.

- **The need for better management and orchestration capabilities:** While ETSI's NFV framework provides for management and orchestration (MANO) capabilities, operators find that existing components and solutions do not live up to their requirements. The emergence of such capabilities, which will allow operators full visibility and control over their networks, would remove the need for the "single neck to choke" which has forced operators to deploy single-vendor stacks.

- **The need for vendors to partner and pre-integrate:** The truth is that standards, blueprints and MANO frameworks are going to take a long time to reach maturity. For this reason, STL Partners believes that vendors must swallow the same "magic pill" as the operators, and commit to making a horizontally-integrated, multi-vendor telco cloud platform a realistic proposition. Vendors are keener to do this than we might expect. We have already begun to see vendors announcing joint propositions. These involve work to pre-integrate both parties' solutions, with the aim of attracting telcos that wish to take a multi-vendor approach but do not have the resource to make it work themselves.

Conclusions: A better telco cloud is possible, and 5G makes it an urgent priority

Telco cloud has not yet delivered on its big promise of openness, flexibility, visibility, control, performance at scale, and offering a platform for agile service innovation.

Yet, it is more important than ever that we figure out how to make telco cloud work – because of 5G. Operators worldwide have committed to investment in 5G radio deployments, which have been justified based on new use-cases which assume extreme sophistication and flexibility in the network.

The operator community is convinced that a multi-vendor telco cloud is the preferred platform to safeguard 5G investments. To make such a platform possible and accessible to the whole industry::

1. Telecoms operators must push harder to ensure that internal transformation initiatives are successful, and that every member of the organisation, from the bottom up, has ownership of delivering telco cloud. This includes rethinking financial processes where appropriate.

2. The industry as a whole must continue to collaborate with the goal of sharing actionable, always-evolving blueprints for telco cloud, rather than rigid standards. This will be of increasing

importance if we are to move to cloud-native, containerised telco cloud environments, which are currently poorly defined.

3. Telecoms vendors must do more to enable the open, multi-vendor ecosystem they have promised, both by pursuing co-opetitive partnerships and pre-integration with other vendors, and by working harder to develop fully-fleshed MANO solutions and products.

If these three things happen, the industry will be on the right trajectory to deliver on the telco cloud promise. Without it, we are unlikely to succeed.

STL Partners is working on a series of research reports which dive into this issue in more detail, aiming to outline how we can ensure telco cloud happens successfully, and quantify what success should look like. Over the coming months, we will publish research on a range of subtopics including containerisation, the move to cloud native, the case for virtualised RAN, and much more.

Table of Contents

Table of Figures

Telco cloud: big promises, undelivered

A network running in the cloud

Back in the early 2010s, the idea that a telecoms operator could run its network in the cloud was earth-shattering. Telecoms networks were complicated and highly-bespoke, and therefore expensive to build, and operate. What if we could find a way to run networks on common, shared resources - like the cloud computing companies do with IT applications? This would be beneficial in a whole host of ways, mostly related to flexibility and efficiency. The industry was sold.

In 2012, ETSI started the ball rolling when it unveiled the Network Functions Virtualisation (NFV) whitepaper, which borrowed the IT world's concept of server-virtualisation and gave it a networking spin. Network functions would cease to be tied to dedicated pieces of equipment, and instead would run inside "virtual machines" (VMs) hosted on generic computing equipment. In essence, network functions would become software apps, known as virtual network functions (VNFs).

Because the software (the VNF) is not tied to hardware, operators would have much more flexibility over how their network is deployed. As long as we figure out a suitable way to control and configure the apps, we should be able to scale deployments up and down to meet requirements at a given time. And as long as we have enough high-volume servers, switches and storage devices connected together, it's as simple as spinning up a new instance of the VNF - much simpler than before, when we needed to procure and deploy dedicated pieces of equipment with hefty price tags attached.

An additional benefit of moving to a software model is that operators have a far greater degree of control than before over where network functions physically reside. NFV infrastructure can directly replace old-school networking equipment in the operator's central offices and points of presence, but the software can in theory run anywhere - in the operator's private centralised data centre, in a datacentre managed by someone else, or even in a public hyperscale cloud. With a bit of re-engineering, it would be possible to distribute resources throughout a network, perhaps placing traffic-intensive user functions in a hub closer to the user, so that less traffic needs to go back and forth to the central control point. The key is that operators are free to choose, and shift workloads around, dependent on what they need to achieve.

The telco cloud promise

Somewhere along the way, we began talking about the *telco cloud*. This is a term that means many things to many people. At its most basic level, it refers specifically to the data centre resources supporting a carrier-grade telecoms network: hardware and software infrastructure, with NFV as the underlying technology. But over time, the term has started to also be associated with *cloud business practices* - that is to say, the innovation-focussed business model of successful cloud computing companies:

Figure 2: Telco cloud defined: New technology and new ways of working

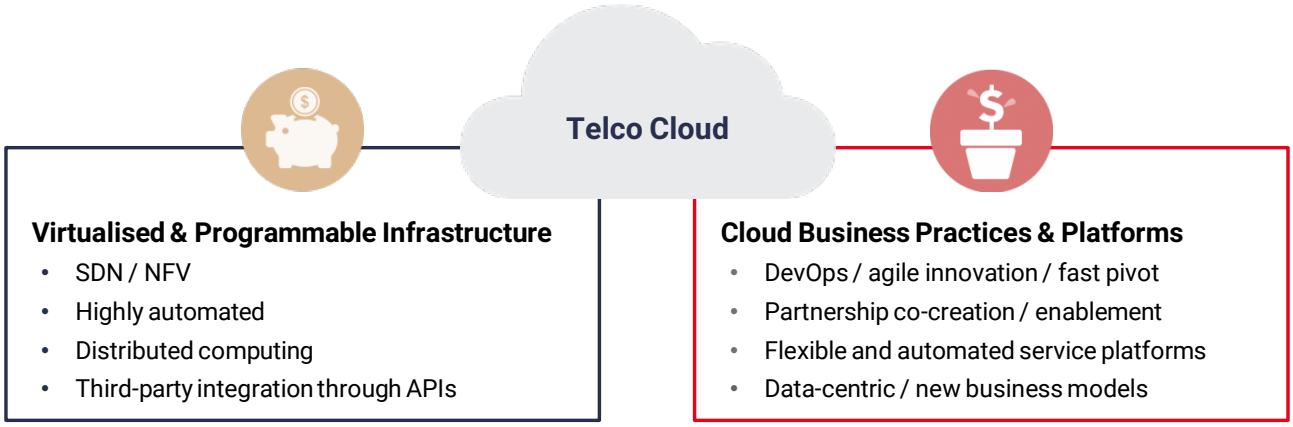

Source: STL Partners

In this model, telco infrastructure becomes a flexible technology platform which can be leveraged to enable new ways of working across an operator's business. Operations become easier to automate. Product development and testing becomes more straightforward - and can happen more quickly than before. With less need for high capital spend on equipment, there is more potential for shorter, success-based funding cycles which promote innovation.

Much has been written about the vast potential of such a telco cloud, by analysts and marketers alike. Indeed, STL Partners has been partial to the same[1]. For this reason, we will avoid a thorough investigation here. Instead, we will use a simplified framework which covers the four major buckets of value which telco cloud is supposed to help us unlock:

Figure 3: The telco cloud promise: Major buckets of value to be unlocked

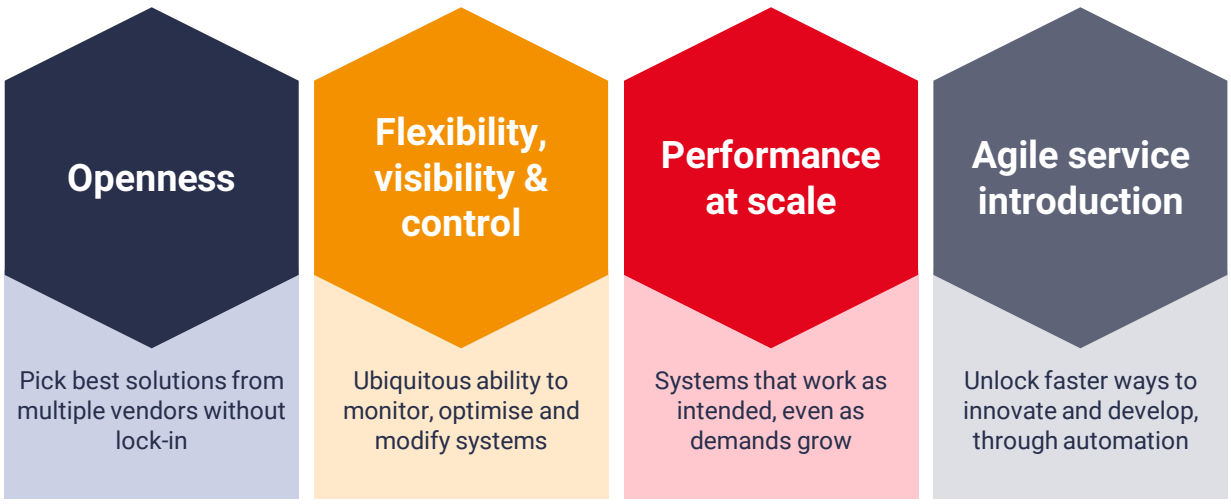

Source: STL Partners

[1] See: Making big beautiful: Multinational operators need the telco cloud

These four buckets cover the most commonly-cited expectations of telcos moving to the cloud. Swallowed within them all, to some extent, is a fifth expectation: cost savings, which have been promised as a side-effect. These expectations have their origin in what the analyst and vendor community has promised - and so, in theory, they should be realistic and achievable.

The less-exciting reality

At STL Partners, we track the progress of telco cloud primarily through our NFV Deployment Tracker[2], a comprehensive database of live deployments of existing telco cloud technologies (NFV, SDN and beyond) in telecoms networks across the planet. The emphasis is on *live* rather than those running in testbeds or as proofs of concept, since we believe this is a fairer reflection of how mature the industry really is in this regard.

What we find is that, after a slow start, telcos have really taken to telco cloud since 2017, where we have seen a surge in deployments:

Figure 4: Total live deployments of telco cloud technology, 2015-2019

Includes NFVi, VNF, SDN deployments running in live production networks, globally[3]

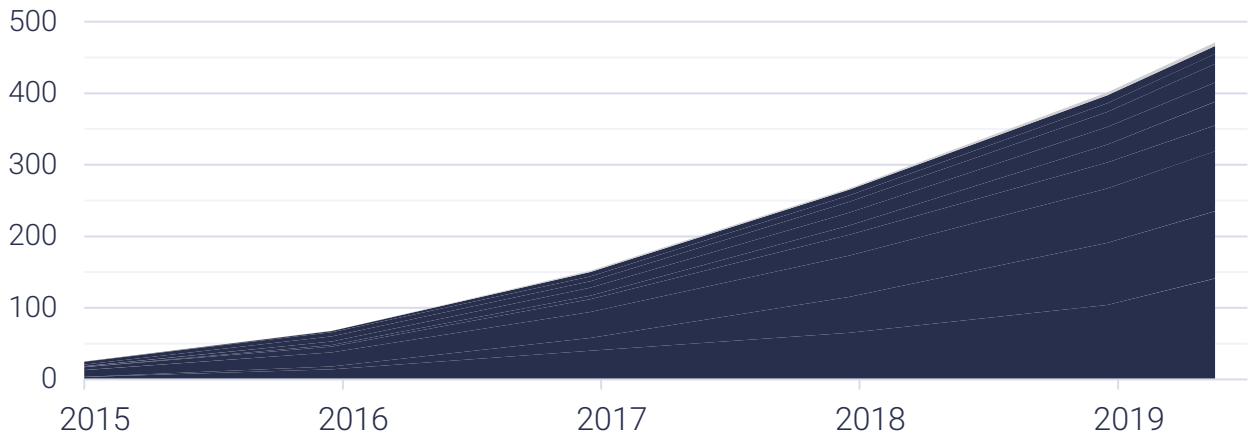

Source: STL Partners NFV Deployment Tracker

All of the major operator groups around the world are now running telco clouds, as well as a significant long tail of smaller players. As we have explained previously[4], the primary driving force in that surge has been the move to virtualise mobile core networks in response to data traffic growth, and in preparation for roll-out of 5G networks. To date, most of it is based on NFV: taking existing physical

[2] See: STL Partners NFV Deployment Tracker
[3] The 500+ deployments here are made up of nearly 1500 total sub-components (network functions, management software, etc). Full definitions and explanations of our counting methodology are given in the NFV tracker itself. Data is a combination of public-domain information and confidential disclosures by service providers.
[4] See: Asia-Pacific points to the future of NFV

core network functions (components of the Evolved Packet Core or the IP Multimedia Subsystem, in most cases) and running them in virtual machines. No operator has completely decommissioned legacy network infrastructure, but in many cases these deployments are already very ambitious, supporting 50% or more of a mobile operator's total network traffic.

Yet, despite a surge in deployments, operators we work with are increasingly frustrated in the results. The technology works, but we are a long way from unlocking the value promised in Figure 2. Solutions to date are far from open and vendor-neutral. The ability to monitor, optimise and modify systems is far from ubiquitous. Performance is acceptable, but nothing to write home about, and not yet proven at mass scale. Examples of truly innovative services built on telco cloud platforms are few and far between.

We are continually asked: will telco cloud really deliver? And what needs to change for that to happen?

The problem: flawed approaches to deployment

Learning from those on the front line

The STL Partners hypothesis is that telco cloud, in and of itself, is not the problem. From a theoretical standpoint, there is no reason that virtualised and programmable network and IT infrastructure cannot be a platform for delivering the telco cloud promise. Instead, we believe that the reason it has not yet delivered is linked to *how* the technology has been deployed, both in terms of the technical architecture, and how the telco has organised itself to operate it.

To test this hypothesis, we conducted primary research with twenty telecoms operators at different stages in their telco cloud journey. We asked them about their deployments to date, how they have been delivered, the challenges encountered, how successful they have been, and how they see things unfolding in the future.

Our sample includes individuals leading telco cloud deployment at a range of mobile, fixed and converged network operators of all shapes and sizes, and in all regions of the world. Titles vary widely, but include Chief Technology Officers, Heads of Technology Exploration and Chief Network Architects. Our criteria were that individuals needed to be knee-deep in their organisation's NFV deployments, not just from a strategic standpoint, but also close to the operational complexities of making it happen.

What we found is that most telco cloud deployments to date fall into two categories, driven by the operator's starting point in making the decision to proceed:

Figure 5: Two starting points for deploying telco cloud

Source: STL Partners

The operators we spoke to were split between these two camps. What we found is that the starting points greatly affect how the technology is deployed. In the coming pages, we will explain both in more detail.

A *function*-first approach to telco cloud

The majority of operators (around three-quarters of our sample) have come at telco cloud because of an immediate need to virtualise a given service or component of their network. In most cases, this is the Evolved Packet Core, which handles data on 4G mobile networks. Usually, traffic growth in mobile

networks is substantial, which means greater capacity is required, and operators wish to support that capacity in a cost-effective, and fairly future-proofed manner, rather than deploying traditional physical infrastructure.

When operators are driven by such needs, there is always an element of urgency. Deployments must be up and running as quickly as possible. For this reason, most operators coming at telco cloud from a function-first perspective have deployed what we term a *vertically-integrated NFV stack*.

To explain what we mean by this, it is helpful to use ETSI's architectural framework for NFV:

Figure 6: ETSI NFV reference architecture framework (simplified)

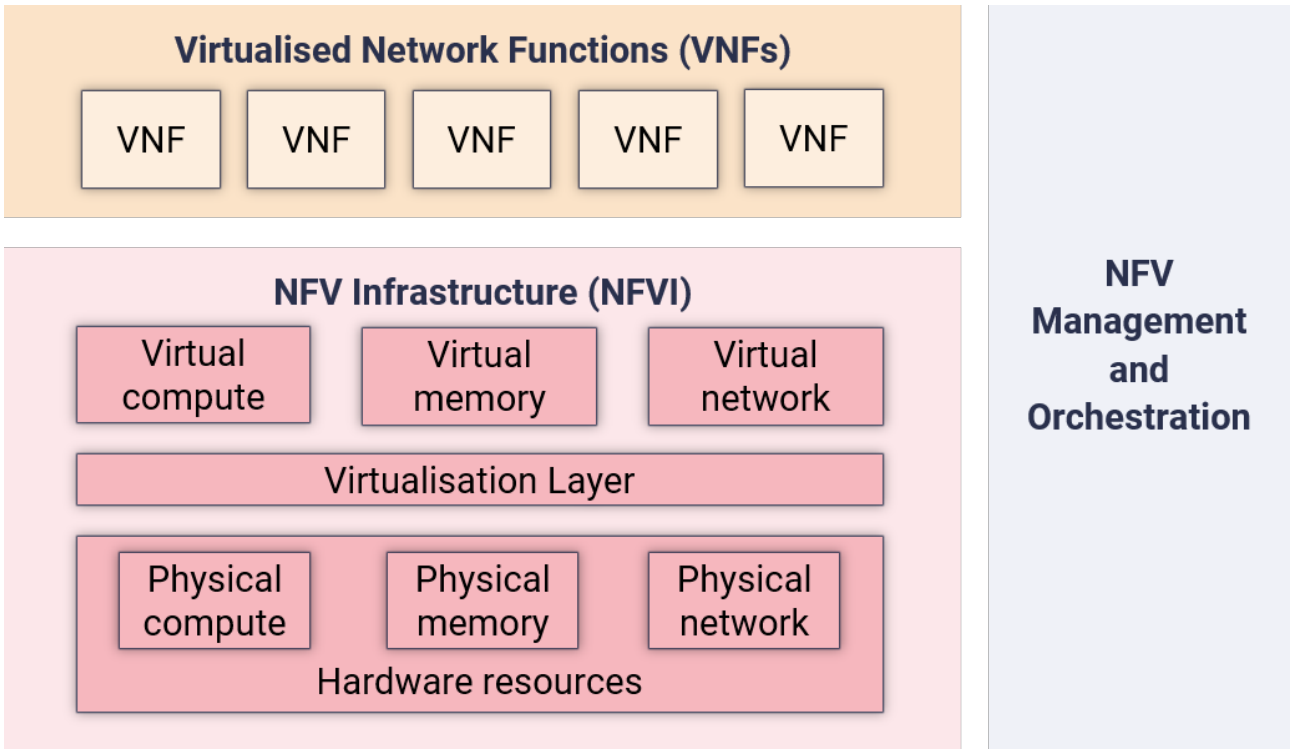

Source: ETSI, via STL Partners

Under this model, network functions run on virtualised infrastructure, with various software components in place to manage and orchestrate the functions and the equipment they run on. In theory, the functions, infrastructure and MANO components are independent of one another, and therefore could be provided by different vendors.

Under the vertically-integrated stack, however, all or most of the components are provided by a single leading vendor partner. In many cases, this is a traditional network equipment provider which may well have worked with the network operator for many years. Examples include Huawei, Ericsson, and Nokia. That partner provides a full-stack solution, including functions, infrastructure and MANO capabilities under one roof that is guaranteed to work. Often, it comes bundled in with the vendor's own computing equipment.

Such an approach is very powerful to an operator with urgent needs to deploy. The vertical stack represents a "one-stop shop", and comes with a comprehensive service-level agreement (SLA) which means that the risk of failure is mitigated. Usually, it plays closely enough to operators' commitments to use open-source software, tending to be based on the vendor's own distribution of the open-source OpenStack platform. It tends to perform well-enough in comparison to old physical infrastructure, when looking at traffic throughput. And, perhaps most importantly, working with existing vendor partners means that it is much more likely to work alongside legacy network infrastructure, which is still far from being decommissioned. In essence, deployment can happen quickly, and fulfil an immediate need.

However, such an approach has several drawbacks, primarily related to working with a single lead partner. Operators complain that while their deployments are based on OpenStack, they are not "open" beyond that. Vertically-integrated stacks of this kind tend to be designed as an integrated package. Operators sometimes (perhaps unfairly) label these packages as "closed ecosystems", citing how difficult they are to integrate with further components or functions when a later decision is made to virtualise something else. For this reason, around half of the operators we spoke to were actually using more than one stack to support different applications. For example, EPC workloads are deployed on one stack, IMS on another, and IT workloads on a third. This means that many operators are running two or more separate telco clouds, which are designed to work for a given purpose, but not to work together. As a result, operators coming at telco cloud from a function-first standpoint have realised very little of the flexibility, visibility and control over their networks that they were promised.

The other problem lies in the way that the VNFs underpinning these deployments were designed. Three to four years ago, VNFs amounted to little more than the old physical functions running inside virtual machines, with minimal effort to optimise for running in a cloud environment. This meant that they would barely outperform the dedicated ASIC-based appliances that were deployed in the past. To get around this, most have deployed hardware acceleration techniques such as *single root input/output virtualisation* (SR-IOV), which transforms a common computing device into one dedicated to a particular function. This has increased throughput - but at the expense of flexibility, because the software and hardware are linked together. To improve performance in the long term, operators want to see network functions rewritten for a cloud environment. This *cloud-native* approach, however, is far from becoming reality.

Overall, we hear that a vertically-integrated approach is successful in the short term, but in the longer term, the cons outweigh the pros:

Figure 7: Pros and cons of the vertically-integrated approach to telco cloud

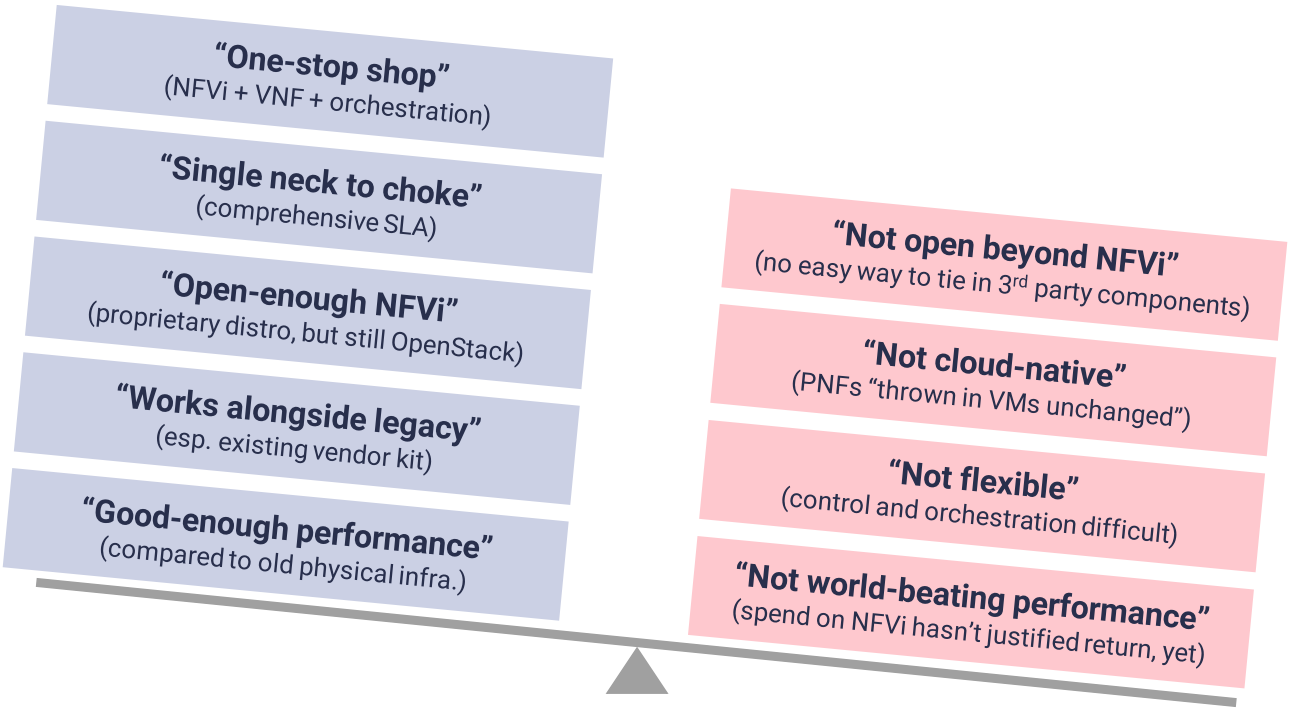

Source: STL Partners, operator research

A *platform*-first approach to telco cloud

If the vertically-integrated stack doesn't deliver, the obvious conclusion might be that we should prioritise building a horizontally-integrated telco cloud which does not suffer from the drawbacks of working with a single vendor. Many of the operators we spoke to agreed.

These operators look to build a single cloud platform that can be used to support any and all telco applications, both now and in the future. Components from different vendors can run side-by-side, in perfect harmony, and swapped in and out at the operator's wish. Through doing so, operators can foster greater competition between vendors, driving performance improvements and ultimately delivering a better end-user experience.

A handful of leading operators have managed to roll out impressive multi-vendor telco clouds, designed from the ground up for their needs. These operators have control over which components are and aren't allowed on the platform, and in all cases have realised significant performance uplift by pitching multiple vendors against each other, forcing them to improve to win.

Of course, it is never that simple. The problem lies in the lack of a practical pre-existing blueprint for a telco cloud platform of this kind. Operators are forced to take a do-it-yourself approach to making it happen – or, as in the case of one Northern European operator, commit to "banging [vendors'] heads together" with contractual obligations to cooperate with one another.

The operators which have succeeded, most of whom are multinational groups, have been able to throw significant resource at defining, building and maintaining their telco cloud platform. Most operators simply do not have the financial or human resource to do this successfully. Even with resource, available, it isn't fast or straightforward: anyone that has followed Telefónica and Vodafone's public announcements related to their NFV programmes over the last few years will be well aware of the many twists, turns and dead-ends they have encountered.

For, while a horizontally-integrated platform is designed to be vendor-agnostic, it does still require the cooperation of vendors to work properly. Our research confirmed that this is an area of concern, even for the leading operators that hold the most influence over the vendor community. In a multi-vendor environment, components of the telco cloud must be designed to work alongside any other component, and there is still much integration work to do to make this happen. Control and orchestration of the network, therefore, are very difficult.

Altogether, this means that while the platform-first, horizontally-integrated approach to building a telco cloud *could* deliver on the telco cloud promise, something needs to change. Again,the cons outweigh the pros for most operators:

Figure 8: Pros and cons of the horizontally-integrated approach to telco cloud

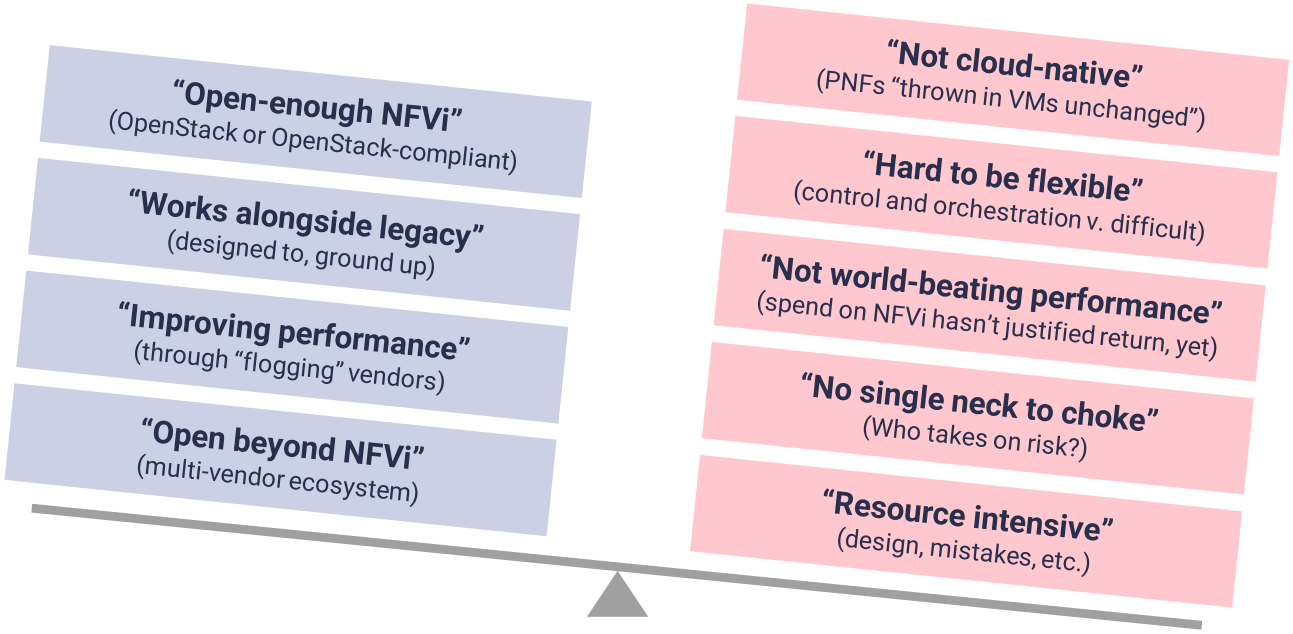

Source: STL Partners, operator research

The solution: change, collaboration and integration

Multi-vendor telco cloud *is* preferred

So far, we have seen two approaches to building a telco cloud, each with benefits, but neither perfect. The function-first approach, which tends to lead to deployment of vertically-integrated stacks, is easy to make happen quickly, but has several drawbacks which mean that it becomes a problem in medium- to long-term. The platform-approach, on the other hand, could be extremely powerful, but requires a great deal of resource and industry cooperation which has so far been lacking:

Figure 9: Pros and cons of function- and platform-first approaches to telco cloud

Source: STL Partners, operator research

From all of our discussions with operators, however, we find that a horizontally-integrated, multi-vendor telco cloud platform *is* seen as the ideal. Those that have come at NFV from a function-first perspective have seen the drawbacks of working with a small pool of vendors. Those who have attempted to build platforms continue to believe that this is the best approach to take to deliver on the telco cloud promise. Consensus is that if the telecoms industry can tackle the negative aspects outlined in Figure 7, we will be well on our way to unlocking the full potential of telco cloud.

In the sections below, we explore how this could happen in practice, from both an internal- and external-facing point of view.

The internal transformation problem

Telco cloud represents a wholesale transformation of a telecoms operator's infrastructure and operations. As such, it requires a wholesale rethink of how the organisation works.

The need to transform is not a new idea, and has been explored in depth by the whole industry over many years, most recently as part of the discussion around "digital transformation". For that reason,

we will not dwell on it here, except to point out some nuances that emerged through the course of our research.

First, most digital transformation frameworks talk in one form or another about the need for whole organisations to "swallow the magic pill" - or commit to change and investing to make it happen. Typically, the recommendation is that those at the top outline a clear vision and roadmap for others to follow. What we have found is that, at least when it comes to networking, top-down direction is not enough. More can and should be done to foster *bottom-up ownership* of building the telco cloud.

Take the example of one Northern European mobile operator, which has an ambitious telco cloud platform strategy. This is pervasive across senior leadership and has led to a variety of significant decisions. One of these is the decision to merge the teams responsible for running networks and those responsible for running IT functions, under a single banner. The rationale is strong: as more and more network workloads move onto common computing infrastructure, there will be increasing need for IT software skills in network operations, and vice versa. Bringing teams together now will facilitate shared learning and development.

In reality, we were told that the new merged group still effectively operates as two siloed entities, to the extent that at lunchtime, IT employees and network employees gather together on different sides of the HQ cafeteria. The reason is simple. While those at the top have articulated a powerful vision, they have not made it relevant to those who actually need to make it happen. The NFV lead at another major European operator cited this as his number one problem, stating that "I am very aware of the need to let go and give my team ownership of our telco cloud, which I find very difficult!".

The extension of this problem is lack of internal skills for operating a telco cloud. As network functions gradually become software, there is increasing need to acquire software development skills which are not part of the core competencies – or at least, self-identity – of a traditional network engineer. Training is expensive, and bringing IT professionals into the mix cannot fully plug the skills gap. Most telcos are looking to hire new blood from outside, either experienced software professionals, or young, cheaper-to-employ graduates with the ability to learn quickly. Both are in high demand, and there are various challenges in attracting them over "sexier" competitors. This is a particular problem in developing markets, where many operators look to systems integrators or their vendor partners to provide the necessary skills. Some telcos have established programmes to tackle this, including vocational partnerships with academic institutions, but it is likely that the real answer lies in more money.

This brings us to the final aspect of the transformation problem: the need to *invest* in network transformation. In all but a few cases, operators we speak to have well-defined mission statements for their telco cloud but tell us that this mission is at odds with the goals of finance departments. Telco cloud involves significant up-front expense, and finance departments demand well-documented proof that the expense will deliver. With very few quantified public proof-points for telco cloud, and with no vendor SLA to guarantee success, the case for spending on a telco cloud platform is weak when

assessed using traditional telco metrics. We have written before[5] about how telcos can rethink financial models and swallow greater amounts of risk. Until this happens, the telco cloud promise is unlikely to be delivered.

The need to foster collaboration and integration

Telcos have been grappling with the need to transform for a long time, with limited success. Our opinion is that, at least in the short-term, the greatest opportunity to accelerate telco cloud lies outside of individual operators' walls.

For a multi-vendor, multi-component telco cloud platform to work, we need to ensure that all the elements can co-exist. We must ensure that partners, vendors, solutions, network functions and frameworks can:

- Co-exist without conflicting with each other
- Co-exist without depending on each other
- Co-exist with legacy networking infrastructure
- Co-exist with legacy IT infrastructure and OSS/BSS

None of this is achievable without significant cross-industry collaboration.

Standards versus blueprints

Our industry has always been fiercely proud of its ability to agree on technical standards for networking technology. After all, standardisation has been key to driving mass uptake of connectivity services across the world. ETSI's NFV Industry Specification Group has been the driving force of such efforts for telco cloud. The NFV framework outlined in Figure 5 is a key part of this. And beyond ETSI, there are many other standards initiatives, including ONAP, OPNFV, MEF, and many more.

Every individual we spoke to was well-versed in the ETSI framework and its alternative, and what it meant to their business. It acts as an important anchor to any discussion about telco cloud. One North American operator cited its adoption of the ONAP framework as a key enabler to its telco cloud platform – and said that without it, they would have made little progress.

But standards are also regarded as a distraction. The truth is that while standards are important, they take time to mature, and in the meantime can only take you so far.

It is helpful to think of telco cloud as an internal-facing product - something built to help fulfil a specific operator's individual needs, rather than to cater for those outside. As such, one operator's telco cloud does not need to be compatible with another's, so long as it serves its internal purpose. A standard

[5] See: Why telco CFOs must start to drive telecoms business model change

might be a useful reference document for building the telco cloud, but it is not necessary to follow it precisely.

That said, there is benefit in collaboration. Without standard frameworks of some kind, VNFs have to be developed to run on multiple custom platforms, which makes it very difficult to productise. Useful testing and validation is near-impossible. As operators experiment with cloud-native approaches to deploying functions, such as containerisation, for which standards are only now beginning to mature, they will need to share learnings and frameworks with others.

As such efforts are emerging to take ETSI's work a step further and define reusable, actionable reference models which take the best aspects of the standards, but leave the operator and vendor partners free to innovate within it. One example is the Common NFV Infrastructure Telecom Taskforce (CNTT), a group formed in early 2019 out of the GSMA and the Linux Foundation.

CNTT's goal is to build "a single overarching reference model [for NFV infrastructure] with the smallest number of Reference Architectures tied to it as possible"[6]. In simple terms, this means a set of documents that outline a more consistent model infrastructure for developers and vendors of telco cloud software and applications to build to, without getting bogged down in slow-moving specifications and restrictive reference architectures.

Figure 10: CNTT's NFVi reference model

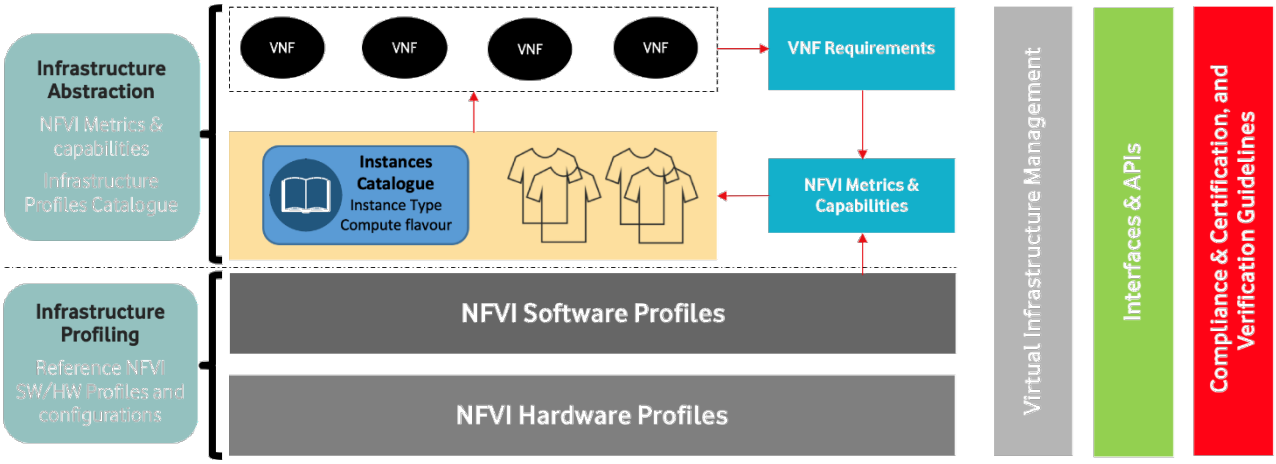

Source: CNTT

It is easy to see the resemblance to Figure 5. CNTT is trying to take the ETSI reference architecture and make it actionable. A large component of this is common interfaces that allow components to co-exist.

The STL Partners view is that CNTT's approach amounts to drawing up an NFV blueprint that can be used and adapted by any operator, giving them confidence that their telco cloud platform will work with whatever components it chooses to deploy. To this end, it is similar to the UNICA programme,

[6] See: CNTT

which outlines a common blueprint for use across Telefónica's operating companies. It is too early to say that the approach will be successful, but we expect to see and hear more about such collaborative initiatives over the next twelve months.

Insufficient management and orchestration solutions

The ETSI NFV framework is supposed to already account for co-existence of different VNFs on a common infrastructure, through its management and orchestration (MANO) component. This contains sub-components intended to manage infrastructure and network functions, bring them together to provision network services, and ensure integration with the operator's back-end systems.

Figure 11: ETSI NFV reference architecture framework (annotated)

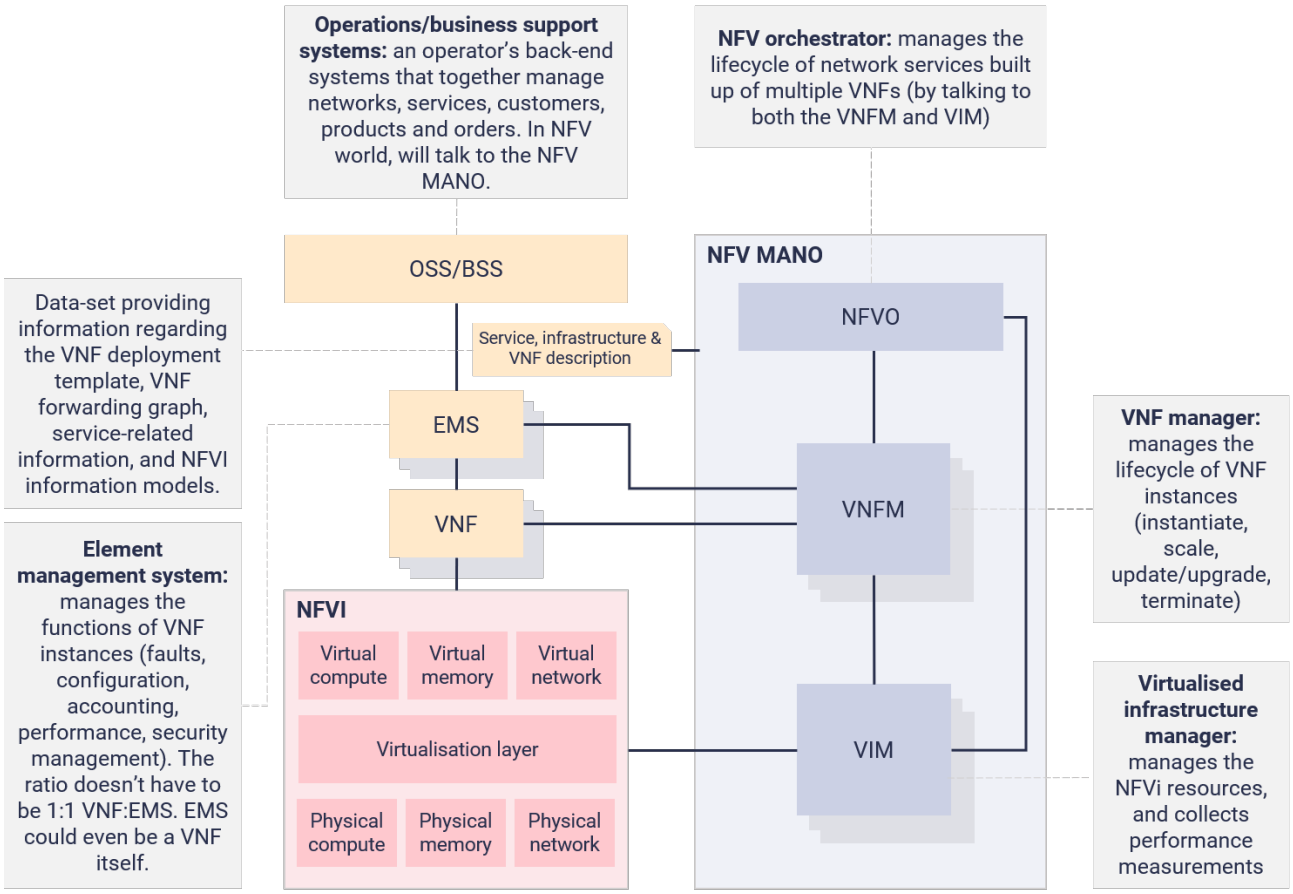

Source: ETSI, via STL Partners

Operators tell us that the market is lacking in software solutions which fulfil their MANO requirements. Individual MANO elements exist and are in use - many vendors sell VNF managers alongside their VNFs, for example. Companies marketing orchestration solutions are easy to find. But, fairly or otherwise, operators believe that nobody can yet offer them a mature end-to-end management and orchestration solution that delivers the flexibility, visibility and control over their network services that they need.

We do not mean to imply that operators are looking for vendors to offer an "uber-MANO" solution that provides all the necessary orchestration functionality as a one-stop shop – nor that this is technically feasible. The key is that operators want the ability to manage and orchestrate infrastructure and network functions regardless of which vendor provided them. This would mean that, when parts of the chain fall over, they would be able to identify where and how, and hold the relevant parties to account. The need for a single neck to choke would disappear.

Vendor partnerships and pre-integration

The truth is that standards, blueprints and MANO frameworks are going to take a long time to reach maturity. They may never do so. Yet telcos want to move forwards as quickly as possible. Deployments will not wait.

For this reason, STL Partners believes that the single greatest opportunity for unlocking the telco cloud promise lies in the vendor community. To survive, vendors must swallow the same "magic pill" as the operators, and commit to making a horizontally-integrated, multi-vendor telco cloud platform a realistic proposition. This means dumping existing business models which focus on building closed ecosystems, and instead embracing a collaborative approach, at all levels of the technology stack.

Vendors are keener to do this than we might expect. We have already begun to see vendors announcing joint propositions. These involve work to pre-integrate both parties' solutions, with the aim of attracting telcos that wish to take a multi-vendor approach but do not have the resource to make it work themselves. Usually, the partnerships involve joint sales and support activity. Examples include Ericsson, which has pre-integrated its VNF and OSS/BSS portfolio with VMware's NFVi platform; and Juniper, which has built a joint solution combining its cloud infrastructure with RedHat's hypervisor.

Some vendors have multiple such partnerships in place. The intention is to acknowledge telcos' concerns that vertically-integrated single vendor solutions do not deliver what they need, and offer something better. We fully expect to see more of this over the coming year.

Conclusion: A better telco cloud is possible, and 5G makes it an urgent priority

We have discussed how telco cloud has not yet delivered on its big promise of openness, flexibility, visibility, control, performance at scale, and offering a platform for agile service innovation. In some cases, this is because operators need to move fast to launch a particular service, and therefore deploy easy-to-action but dead-end vertically-integrated solutions. In others, operators look to build horizontally-integrated, multi-vendor telco cloud platforms which will cater to all their future needs, and find that this requires high levels of resource. The industry is, therefore, in somewhat of a quandary.

Yet, it is more important than ever that we figure out how to make telco cloud work – because of 5G. Operators worldwide have committed to investment in 5G radio deployments, which have been justified based on new use-cases which assume extreme sophistication and flexibility in the network. Network slicing, ultra-low latency applications, and so on – all of these assume end-to-end virtualised infrastructure which does not yet exist.

The operator community is convinced that a multi-vendor telco cloud is the preferred platform to enable and safeguard investments in 5G. Hence it is imperative that we work together to make such a platform possible and accessible to the whole industry:

1. Telecoms operators must push harder to ensure that internal transformation initiatives are successful, and that every member of the organisation, from the bottom up, has ownership of delivering telco cloud. This includes rethinking financial processes where appropriate.

2. The industry as a whole must continue to collaborate with the goal of sharing actionable, always-evolving blueprints for telco cloud, rather than rigid standards. This will be of increasing importance if we are to move to cloud-native, containerised telco cloud environments, which are currently poorly defined.

3. Telecoms vendors must do more to enable the open, multi-vendor ecosystem they have promised, both by pursuing co-opetitive partnerships and pre-integration with other vendors, and by working harder to develop fully fleshed MANO solutions and products.

If these three things happen, the industry will be on the right trajectory to deliver on the 5G promise. Without it, we are unlikely to succeed.

STL Partners is working on a series of research reports which dive into this issue in more detail, aiming to outline how we can ensure telco cloud happens successfully, and quantify what success should look like. Over the coming months, we will publish research on a range of subtopics including containerisation, the move to cloud native, the case for virtualised RAN, and much more. As always, our work will be grounded in what is happening on the ground in telecoms operators. We welcome your feedback.

Index

STL PARTNERS

Research

Consulting

Events

www.stlpartners.com

PN3200081-001-EN