Vendor-hyperscaler partnerships: Time to bring in the telcos

Telcos have tended to view hyperscalers with suspicion

It would be fair to say that telcos' implementation of cloud-native network functions (CNFs: including the 5G SA core and open/vRAN) and of vendor-neutral, horizontal telco cloud platforms has taken longer than we initially expected. Telcos have been deliberating: should they build their own, private telco cloud? Should they, instead, opt for deployment on the hyperscale cloud as a faster and less expensive means to achieve the same goal? Or should they mix and match, i.e. go hybrid?

Much of this has been due to telcos' initial reticence about working with hyperscalers. First, a bit of background. A few years ago, telcos were scared that hyperscalers would muscle into the connectivity business and 'eat their lunch'. The likes of AWS and Microsoft Azure were busy retrofitting their clouds to carry network workloads. A big step forward in this process came in 2020 and 2021, when Azure bought two leading challenger vendors and exponents of cloud-native network functions (CNFs) – Metaswitch and Affirmed Networks – followed by its acquisition of AT&T's entire Network Cloud (= private telco cloud) platform. Along with the technology, all three deals brought networking expertise from all three acquired entities, or 'telco DNA' as Azure calls it.

Did this mean that hyperscalers were going to play at being both telco and vendor? And would hyperscalers leverage the telco expertise and cloud-native networking capabilities they were rapidly building up – through acquisitions of and partnerships with vendors – to bypass telcos, and deliver edge computing and other network-dependent services (such as IoT, Industry 4.0 and the like) direct to their enterprise clients? Let alone regular, commodity connectivity?

Telcos' wariness towards collaborating with hyperscalers extended to the question of how it would work in practice to deploy Virtual Network Functions (VNFs) and CNFs, and run the associated workloads, on the public cloud. Was the technology mature and sufficiently 'carrier-grade'; what operational changes would be needed to run network functions on public cloud infrastructure; and would reliance on public cloud leave telcos prey to a double form of lock-in: both to the vendors that the hyperscaler supported, and to the hyperscaler itself – its platform and roadmap?

Hyperscalers and vendors have collaborated to enable CNF deployment on public and hybrid cloud

While telcos pondered, collaborations were gathering pace towards adapting and testing vendors' core and RAN CNFs on the big-three hyperscale clouds (including Google Cloud). In our latest report, Vendor-hyperscaler partnerships: Building the multi-cloud native future?, we look at the R&D, PoCs and commercial deployments involving the big three hyperscalers, on the one hand, and the vendors Ericsson, Mavenir, Nokia and Samsung, on the other.

We have to admit that we ourselves have been somewhat surprised at the depth, scope and maturity of these collaborations. In addition, vendors have made a lot of progress in adapting their CNFs to run on multiple container orchestration (aka Kubernetes) platforms over both private and public (and by extension, hybrid) cloud. The (self-proclaimed) readiness of all four of these vendors' CNFs – RAN, core or both – to run on

multiple Kubernetes variants, including the hyperscalers', is illustrated below. We describe and analyse these collaborations in detail in the report.

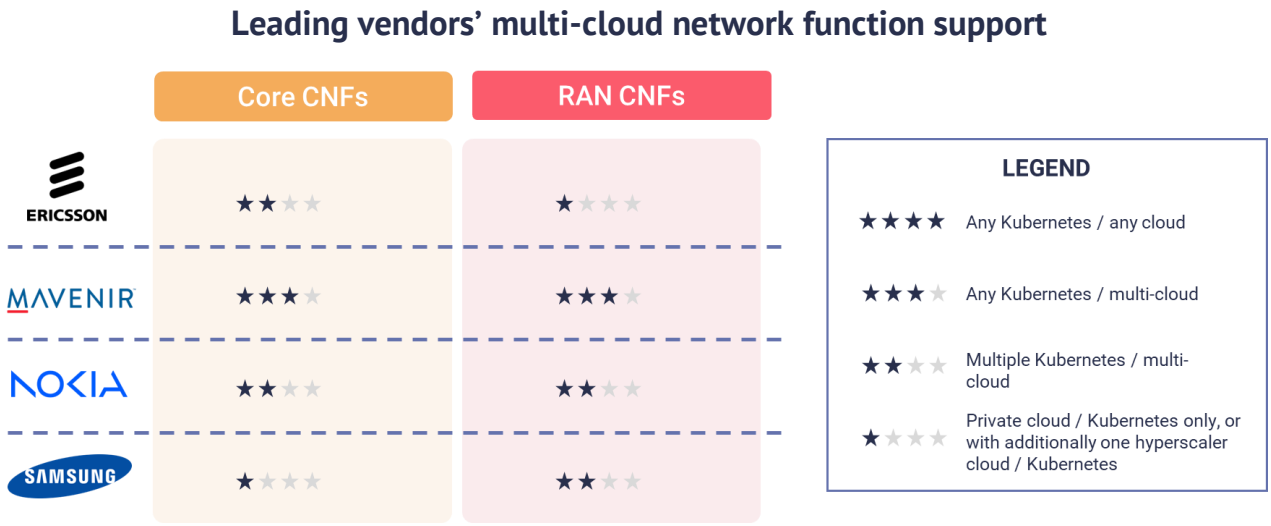

Source: STL Partners

Scant deployments to date but new announcements this year

Despite the vendors' and hyperscalers' claims that their combined platforms are ready for scaled, commercial deployment, very few such deployments have yet taken place or got underway. Indeed, up until this year, the only examples of deployments supporting commercial, macro-network services were:

- DISH Network's deployment on the AWS cloud of the core and some of the RAN functions supporting its greenfield 5G network in the US (still ongoing)
- And AT&T's (still ongoing) deployment of its 5G SA core on Microsoft Azure Operator Nexus: the commercial version of the former AT&T Network Cloud that has been fully integrated and remoulded into a Microsoft product available for third parties, in addition to AT&T itself.

However, this year, there have been two further announcements of significant deployments on public cloud, with the spoils shared equally between Azure and AWS:

- In February, e& UAE announced it would deploy the Azure Operator 5G Core – the Microsoft iteration of the Affirmed Networks core technology – on the Azure Operator Nexus platform running in e& facilities.
- And in May, Telefonica O2 Germany announced that it would deploy a Nokia 5G SA core on the AWS cloud – although details about which physical facilities the cloud platform would be running on were not available at that stage.

Are we at the start of scaled adoption?

It is encouraging that DISH and AT&T have been joined by two further telco cloud innovators in deploying CNFs on public and hybrid cloud. In three out of these four deployments, the CNFs at stake are limited to the 5G SA core. There have been considerable doubts in the industry up to now about whether the hyperscale cloud is sufficiently telco-grade to fulfil the performance requirements of SA data plane workloads, let alone RAN workloads. Depending on how they work out in respect of timescales, capex and opex, performance, and

user experience, these deployments may begin to inspire confidence in the operator community more generally about the technological maturity of leading vendors' core CNFs to be run on the public cloud.

Indeed, one of the reasons why roll-outs of the 5G SA core as a whole have taken much longer than originally expected is that operators were unsure of the degree to which they should deploy these – in part or whole – on the hyperscale cloud. It is possible, then, that success in these initial public cloud-based roll-outs will help catalyse the implementation of 5G SA by other telcos, with some of the new deployments also carried out on public – or at least hybrid – cloud.

Deployments will be increasingly hybrid and multi-cloud

At least three of the four deployments described above are also hybrid, i.e. being deployed on the hyperscaler's cloud platform but also in part or whole on the telco's own physical facilities: AT&T, DISH and e&. And, as we observed above, it has not yet been disclosed whether O2 Germany's roll-out over AWS will follow the same pattern.

The role of hybrid cloud will be even more important in the case of deployments of the virtualised and/or open Radio Access Network (vRAN and open RAN). This is because RAN has stringent requirements in terms of real-time data processing at high volumes, particularly for the Distributed Unit (DU) function: part of the baseband function in mobile networks, sited typically at the foot of the cell tower or close to it. Standard hyperscaler compute platforms are inadequate to this task; so, workload-specific hardware acceleration and custom silicon will be needed, which will be run on telco facilities even if the Central Unit (CU) of the baseband is hosted on the public cloud.

Such hybrid RAN platforms are now being offered commercially by Azure for the private network market (in partnership with Nokia), and by DISH and AWS as part of DISH's macro-network roll-out (with Mavenir, Rakuten Symphony and Samsung as the RAN software vendors).[1]

The tech case is made: now make the business case work for telcos

Let us assume that the scaled deployments above prove the technical case for deploying telco workloads on hybrid cloud. But what is the business case for telcos of scaled migration? In addition to potential benefits of cost reduction, scalability and flexibility, we believe that the ability for telcos to create and operate network services across multi-cloud infrastructure will be key for them to capture the opportunities from an emerging networking paradigm that we have previously called Network-as-a-service (NaaS) 2.0.

This involves the evolution of networking into a software-based and cloud-delivered, on-demand service based around the needs of vertical-specific, real-world use cases and processes – with AI and gen AI helping to power the optimisation and integration of networking to deliver the desired outcomes.

[1] Another example of a hybrid deployment of this type is NTT DoCoMo's deployment of AWS's distribution of Kubernetes at DoCoMo's RAN sites to support flexible switching of RAN workloads between DoCoMo's private cloud infrastructure and the AWS cloud, for the purposes of resilience, surges in user demand and disaster recovery. In other words, as back-up.

Complementary roles for vendors, hyperscalers and telcos in application networking

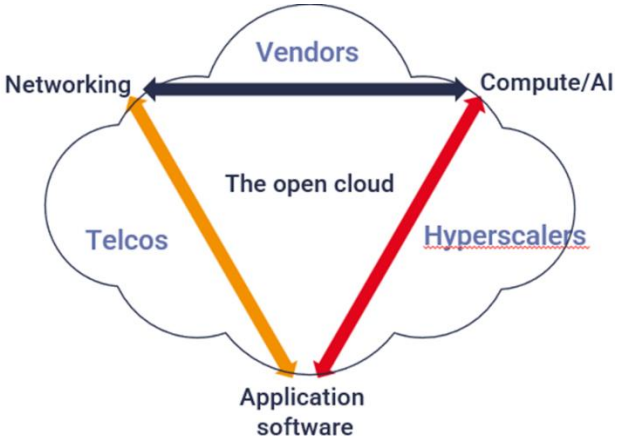

Source: STL Partners

The diagram illustrates how vendors, hyperscalers and telcos (and other networking providers) can contribute their respective strengths and assets to an ecosystem supporting this new 'application-networking' economy. We describe the complementary roles of vendors, hyperscalers and operators in detail in our report.

In essence, in the application-networking economy: **hyperscalers provide the infrastructure and the platform, while network service providers (telcos and others) operationalise, commercialise and deliver those services in the physical world, via both virtual and physical network technology developed by vendors.**

Telcos need to have confidence that they can generate significant value from a role in application networking / NaaS 2.0 such as this. But for this to come about, there needs to be a shared vision and purpose on the part of vendors and hyperscalers, too. All three have complementary strengths and assets that they can bring to the deal.