

Operators' cloud strategies for vRAN are challenging, but a USD23.5 billion far edge market awaits

January 2021

Caroline Chappell and Caroline Gabriel

Operators want to run RAN software on a cloud platform in order to gain the agility and cost benefits of cloud-native automation, commodity hardware and resource flexibility and scalability. Operators with a horizontal, 'open' RAN cloud strategy further wish to select best-of-breed RAN functions rather than relying on a single vendor to deliver a monolithic network. However, 5G RAN is a highly complex network domain with stringent performance and latency demands and huge processing requirements. Implementing a cloud platform for the virtualised RAN (vRAN) will involve tearing up the rule book for the centralised IT clouds that operators have built so far to support other network functions. Projects to cloudify the RAN are at risk of failure unless operators and vendors can build the new, distributed, high-performance cloud platform that is needed to support the 5G vRAN. This difficult and expensive exercise may be mitigated if the vRAN cloud provides the right environment for running an emerging breed of enterprise use cases with similar latency and performance requirements to vRAN functions.

Virtualisation of the RAN demands a new cloud platform

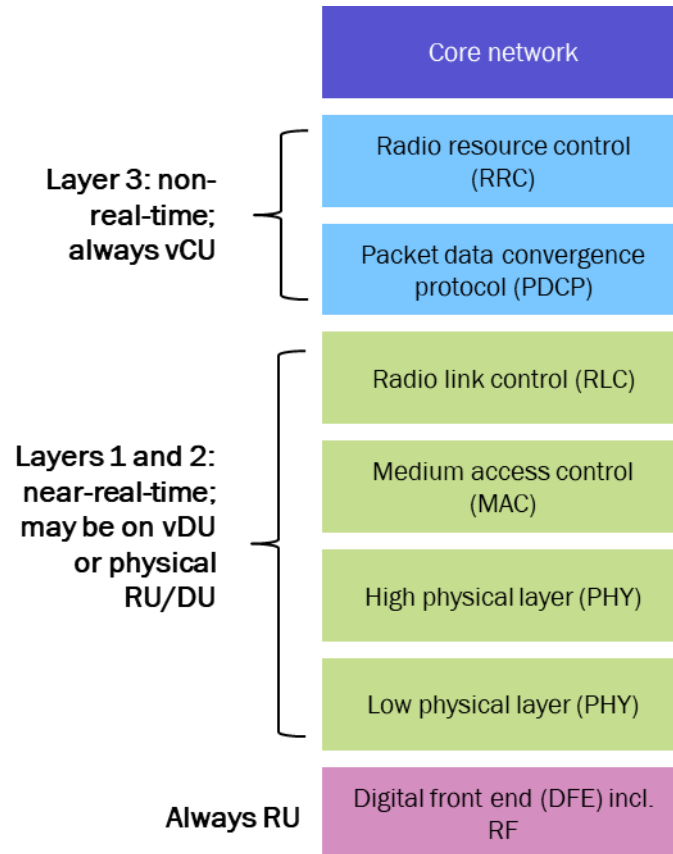
The RAN is the last major network domain to be virtualised. This involves disaggregating RAN software from the highly engineered, proprietary hardware and chipsets that are only available from the handful of specialist vendors that have supported the RAN to date. The arguments for applying cloud technology to the RAN are similar to those for other domains; that is, to give software access to a lower-cost and more agile cloud platform. Open RAN initiatives point to a further incentive for enabling RAN software to run on a vendor-neutral, horizontal cloud platform: with the appropriate standardisation and APIs, different RAN functions can be procured from multiple vendors, thereby breaking the monopoly of traditional vendors, encouraging innovation and driving down function prices.

However, the RAN differs from other network functions in several key ways, as described in our report, [*Implementing the vRAN cloud: strategies for success*](#). As such, it makes new demands on the network cloud.

- The RAN cannot be fully virtualised. Six functions are amenable to running on a vRAN cloud (see Figure 1), but significant components such as radio units (RUs) remain physical, and the digital front end (DFE; including RF functions) remain embedded in dedicated hardware.
- Layer 1 and 2 functions (using the vDU) must be as close as possible to cell sites and incoming signals if they are to support 5G-standard response times, so RAN functions need to be distributed on a scale well beyond that required for other network functions. A vRAN cloud will need to run as a single logical platform across thousands of cell sites, whereas operator network clouds today are typically discrete instances running in a handful of centralised data centres.
- 5G RAN functions have highly demanding processing requirements due to increased bandwidth (100MHz channels instead of 20MHz for 4G), an increased number of antenna branches (massive MIMO requires 20 times the baseband processing power of 4G) and the need for 0.2ms transmission times compared to 3ms in

LTE. An x86-based cloud might deliver only 15–20% of the performance of a dedicated baseband platform for low-layer 5G functions and will need hardware acceleration rather than the software-based acceleration that is implemented in network clouds today. Efficient support for heterogeneous hardware acceleration based on GPUs, FPGAs and/or ASICs remains an unresolved issue.

Figure 1: Simplified RAN stack showing the six CU/DU functions that can be virtualised



Source: Analysys Mason, 2021

Network clouds based on current network functions virtualisation (NFV) technologies will not support the lower-layer functions of a vRAN. Emerging cloud-native platforms are better suited for resource-constrained, low-latency vDU environments, but they are IT-oriented and immature. They require significant development in areas such as security, IP addressing, hardware acceleration abstraction and bare metal support before they can be used for vRAN. Public cloud and established enterprise cloud technology stack providers do not yet address these issues and operators are still finding their way with first-generation, NFV clouds. The vRAN project may stall or even fail unless vRAN cloud challenges can be resolved.


The vRAN cloud is a blueprint for the edge cloud that is needed to support demanding industrial use cases in the future

The telecoms industry has a near-term and well-scoped use case (the vRAN cloud) for a highly distributed, low-latency, high-performance cloud platform and is leading the way in its development. However, the vRAN cloud's stringent requirements mean that it will be expensive to build. Operators expect that virtualisation will deliver total cost of ownership benefits over the RAN's lifetime, but the cost reduction in the first few years of a

deployment will not be sufficient to justify the investment on its own. Operators need to justify a vRAN cloud partly on the grounds of the new revenue streams that it can enable.

Use cases associated with the transformation of industrial operations have similar latency/location/hardware acceleration requirements to vRAN functions. Such use cases rely on data-processing-hungry technologies such as AR/VR, holography, robotics/cobotics and AI/ML-driven computer vision/natural language/telemetry data processing, which are expected to revolutionise the speed of business decision making and the quality of customer experience. Enterprises expect to eventually run such use cases across hundreds of thousands of micro data centre locations, which most will be unable or unwilling to build themselves. Operator vRAN clouds will provide both the scale and the properties needed by the applications that underpin industrial use cases. We estimate that cloud compute at the far edge will generate USD23.5 billion of revenue by 2030(see Figure 2), out of an overall public edge market worth USD70 billion. Operators will be well-placed to win a large share of such revenue if they have deployed vRAN clouds.

Figure 2: Enterprise spending on public edge services (IaaS, PaaS, SaaS) at the far edge, and percentage of the total public edge spending, worldwide, 2019–2030



Source: Analysys Mason, 2021

Operators that are planning to adopt an open, multi-vendor edge cloud that will run disaggregated RAN functions will need to push the market to create a new technology stack that is capable of supporting vRAN processing and performance requirements. Operators should collaborate with the right partners to take advantage of their need for cloud at the far edge and accelerate the development of advanced enterprise use cases that can benefit from such clouds. This will secure their position in the emerging edge value chain and help with the business case for the vRAN cloud.