



ALTIOSTAR
Leading Network Transformation



Perspective

An automated, platform-based approach to RAN virtualization is key to achieving its benefits

November 2021

Caroline Chappell

Contents

1.	Executive summary	1
1.1	Automation is a key component of the business case for virtualized RAN	1
1.2	A cloud platform approach to automation mitigates risk	2
1.3	Building zero-touch automation for the virtualized RAN	2
2.	Virtualizing the RAN creates new opportunities to reduce TCO and generate revenue	3
2.1	What is a virtualized RAN?	3
2.2	The benefits of an open, disaggregated approach to the RAN	5
2.3	Risks of deploying a disaggregated RAN	6
3.	A platform architecture for structuring an open, virtualized RAN mitigates disaggregation risks	7
3.1	What is a platform?	7
3.2	Platform-based automation mitigates disaggregation risks	9
4.	The benefits of cloud-native portability and zero-touch automation for a virtualized RAN	10
4.1	Designing automation for the virtualized RAN	10
4.2	Benefits of automating a virtualized RAN	11
4.3	Selecting a platform for zero-touch automation	12
5.	Conclusion	12
6.	About the author	13

List of figures

Figure 1.1: Three technology pillars supporting a virtualized RAN.....	1
Figure 2.1: Overview of the architecture for the legacy RAN, the one-split vRAN and the disaggregated vRAN.....	4
Figure 2.2: Commercial drivers for deploying vRAN before 2026	6
Figure 3.1: Three layers of platform functionality.....	8
Figure 3.2: A vRAN architected as a platform	9

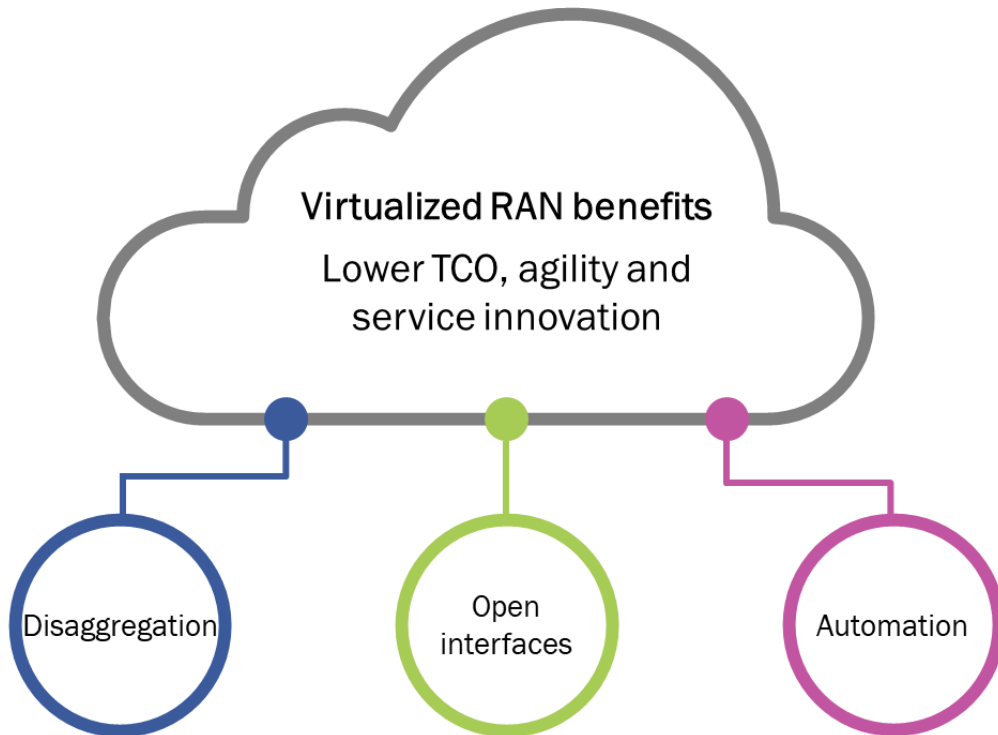
1. Executive summary

1.1 Automation is a key component of the business case for virtualized RAN

Mobile network operators (MNOs) worldwide have been slow to virtualize the radio access network (RAN) because of its stringent performance requirements and complexity. However, they are now turning their attention to this expensive and demanding area of the network in search of the types of benefit that telecoms operators are beginning to experience in other network domains, including lower total cost of ownership (TCO), greater agility in scaling capacity and the prospect of increased and faster service innovation in future.

RAN virtualization is being enabled by the confluence of emerging open standards that support its functional disaggregation (into central, distributed and radio units) with new technologies that facilitate its technical disaggregation (separating RAN software from proprietary hardware). This new, open, interface-based approach to RAN virtualization is now possible because cloud-native technologies are more amenable to automation than the previous set of technologies used for network function virtualization (NFV). Software for the virtualized RAN (vRAN) is currently being developed – or re-developed – in a cloud-native way as microservices destined for containers running on commodity, cloud-native infrastructure. Automation (orchestration) is an inseparable consequence of using modern software design and cloud-native technologies to develop and deploy virtualized RAN functions. Zero-touch automation, using standard cloud-native tooling and approaches, is key to achieving the benefits that MNOs expect from a virtualized RAN, particularly improved agility and innovation, although automation’s role in reducing opex is also important. Together, technical disaggregation, open interfaces and automation make the business case for RAN virtualization (see Figure 1.1).

Figure 1.1: Three technology pillars supporting a virtualized RAN



Source: Analysys Mason

1.2 A cloud platform approach to automation mitigates risk

Analysys Mason expects MNOs to invest USD12.1 billion in cloud infrastructure by 2026 to support their vRANs. Building a vRAN using a cloud-native platform approach addresses the risks of disaggregation and enables MNOs/integrators to develop management and automation in a consistent and reusable way across vRAN components. In a cloud-native, virtualized RAN context, the cloud infrastructure and management and automation layers of the platform are based on open and open-source tooling being developed within the Kubernetes ecosystem, under the auspices of the Cloud-Native Computing Foundation (CNCF).¹ It is a striking feature of a cloud-native virtualized RAN that this tooling supports the automation of all its software components, from the cloud infrastructure itself all the way up to the cloud-native network functions, such as the central and distributed units (CUs and DUs).

A cloud-native platform extends the Kubernetes ecosystem tooling by applying artificial intelligence (AI) and machine learning (ML) and analytics capabilities: advanced MNOs that already have leading-edge deployments of a virtualized RAN point to the importance of having a data and AI/analytics layer that is shared by all components in the vRAN solution as this provides a single source of truth for virtualized RAN operation and management.

Use of such a cloud-native platform lowers the risks associated with virtualizing and disaggregating a network domain as large and complex as the RAN. A platform approach, with common tooling, shared data, open interfaces and an ability to enforce automation consistency across all components, reduces integration efforts. It provides a common environment in which software products from multiple vendors can plug and play interchangeably, mitigating the risk of lock-in. Zero-touch automation, built using shared platform tooling (including AI/analytics tools) and the scaling and self-healing capabilities of a cloud-native platform, addresses the challenges of managing such a highly distributed system as the RAN at scale. This approach provides greater reliability and resilience than is possible with manual operations. The fact that such a cloud-native platform underpins business IT systems, as well as network functions, enables unprecedented integration between business and network processes. MNOs will be able to use business intelligence to drive automated decisions at network level, resulting in a virtualized RAN that is optimized for the business.

1.3 Building zero-touch automation for the virtualized RAN

Zero-touch automation needs to span three phases of a virtualized RAN lifecycle: Day 0 deployment, Day 1 preparation for operation and Day 2 maintenance for the period in which the vRAN is live in production.² MNOs need to build their RAN lifecycle management automation in conjunction with their RAN vendors, whose automation playbooks will be needed to run the software components on the common platform. MNOs need to accommodate the iterative nature of automation development: it will take time to evolve automation towards zero-touch, AI-driven control that can predict anomalous events and make appropriate changes to the virtualized RAN to counter them. MNOs should also develop their RAN automation in a modular way and at a suitable level of abstraction so that cloud-native tooling and approaches can be used across additional network domains in future, further reducing the cost and complexity of the network end-to-end.

This paper points out that the provider of the cloud platform on which a virtualized and open RAN will run will play a critical role in zero-touch automation. The cloud platform provider will be the source of the tooling on which such automation will be built, so operators are advised to select their RAN platform carefully. They should pay attention to the cloud platform provider's level of support for industry-leading, open-source DevOps

¹ For more information about the CNCF, please see CNCF's *Who we are*. Available at: <https://www.cncf.io/about/who-we-are/>.

² See Section 4.1 of this perspective for further details about these three phases.

tooling, its track record of working with MNOs and network function vendors on network virtualization programmes and its experience and expertise around cloud-native automation.

2. Virtualizing the RAN creates new opportunities to reduce TCO and generate revenue

2.1 What is a virtualized RAN?

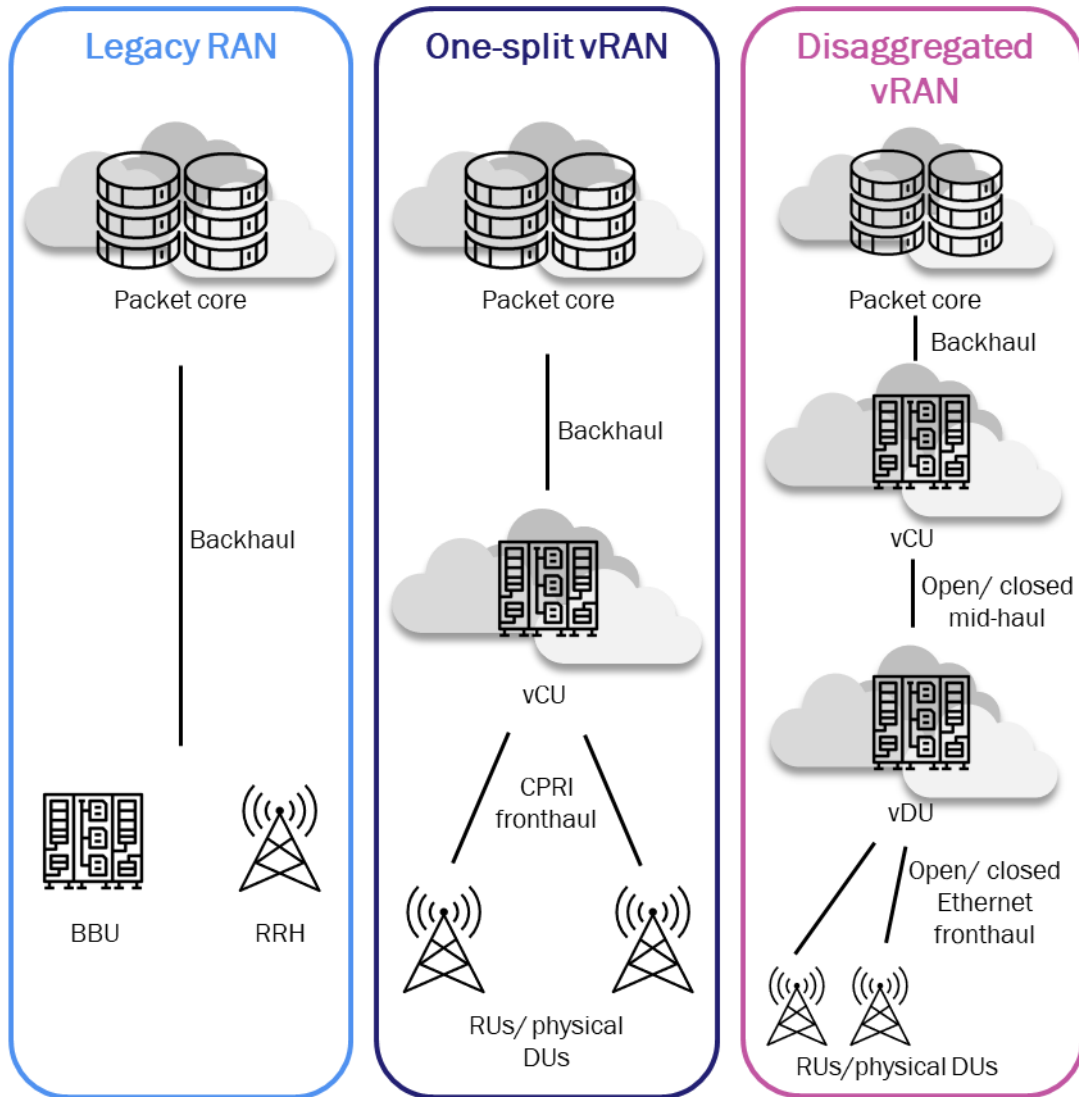
The RAN is a highly complex and critical network. Until recently, it resisted the process of virtualization seen in other network domains because of its significant performance requirements. In fact, the RAN cannot be entirely virtualized, unlike many other network elements. Radio units (RUs) will remain physical and will continue to be embedded in dedicated hardware. However, non-RU functions can be virtualized. Today, as MNOs plan their RAN virtualization roadmaps, vRAN will be the fastest-growing domain in terms of cloud investment: it will grow from USD293 million in 2021 to USD12.1 billion by 2026.³ According to recent Analysys Mason research, most operators want to implement a virtualized RAN with open interfaces, that is, an Open vRAN. We will refer to the Open vRAN as ‘vRAN’ in this paper for simplicity.

There are multiple options for realizing a vRAN architecture, as Figure 2.1 shows. The **one-split vRAN** is the simplest option: high-layer, non-real-time RAN functions are virtualized and run in a central unit (CU) in the cloud, while the near-real-time, lower-layer RAN functions are run in a physical distributed unit (DU), which forms a single logical unit with the RUs.

Open vRAN initiatives favour a **disaggregated vRAN**, in which some or all lower-layer functions are abstracted and run on a virtualized DU. This is logically separate from the physical RUs. Open vRAN initiatives mandate open, standardized interfaces (APIs) between the CU/DUs/RUs but a disaggregated vRAN can still be realized using vendor-proprietary interfaces.

³ For more information, see Analysys Mason’s *Network cloud infrastructure: worldwide forecast 2021–2026*. Available at: <https://www.analysismason.com/research/content/reports/cloud-infrastructure-forecast-rma16/>.

Figure 2.1: Overview of the architecture for the legacy RAN, the one-split vRAN and the disaggregated vRAN



Source: Analysys Mason

The important point to note here is that the functional disaggregation of vRAN components is facilitated by the technical disaggregation of vRAN software from proprietary hardware as part of a network virtualization process. Virtualization technologies have evolved over the past decade, and vRAN is taking advantage of the next cloud-native generation of technologies. 10 years ago, network function virtualization (NFV) involved extracting monolithic network function software from proprietary hardware and running it in virtual machines on commercial off-the-shelf (COTS) servers, that is, in ‘virtual appliances’. Software for the vRAN is today being developed, or re-developed, in a cloud-native way, as microservices destined for containers running on commodity, cloud-native infrastructure. In order to support the virtualized RAN, whether as an open vRAN or one from a single vendor, MNOs will need to understand the role of the cloud-native platform as the management and execution environment for all vRAN software components, and the high level of automation such a platform enables.

This paper will focus on the automation of a virtualized, disaggregated RAN with open interfaces between vRAN functions. Such open interfaces enable different vRAN functions to be supplied by different/multiple vendors.

2.2 The benefits of an open, disaggregated approach to the RAN

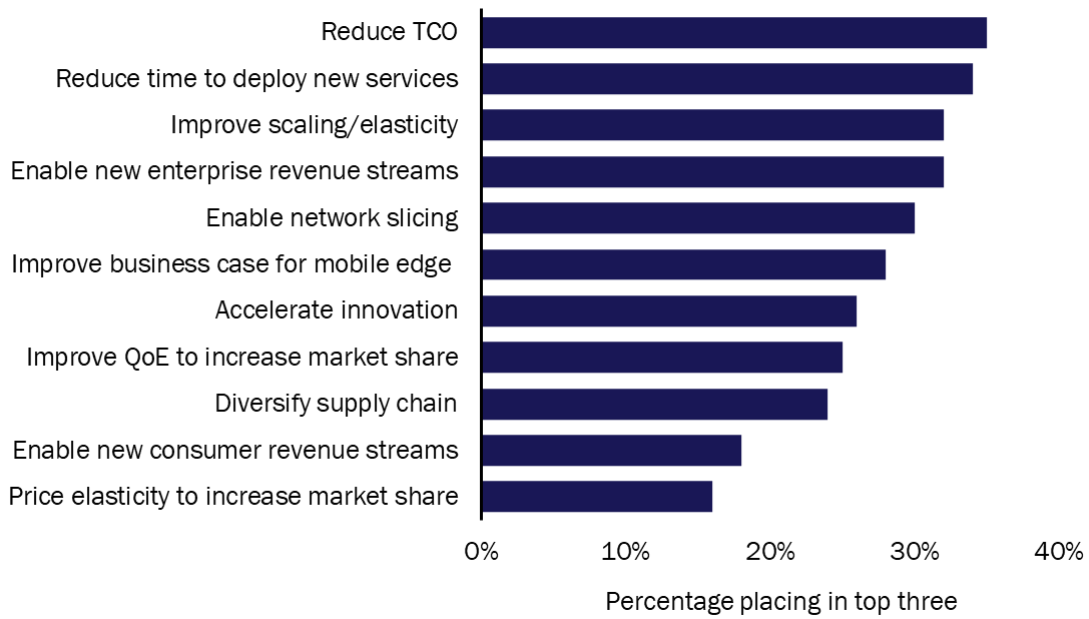
MNOs expect to obtain an increasingly diverse set of benefits from the disaggregation and virtualization of the RAN. Analysys Mason's survey of 74 mobile network providers (conducted in 3Q 2020 and supported by other research) indicates that total cost of ownership is losing ground as a driver for RAN virtualization. Only 35% of respondents cited it as a top-3 driver in 2020, compared to 65% in 2018, and it came out only marginally ahead of 'reduce time to deploy new services' (34% of respondents). From previous experience NFV, MNOs understand that they can gain capex benefits from using lower-cost commodity components, such as servers and open-source software and they also expect that greater competition comes with a multi-vendor landscape to drive down costs in a vRAN. However, opex benefits from NFV have been harder to attain due to the complexity of NFV orchestration.

This situation changes in a cloud-native vRAN environment. Automation (orchestration) is an inseparable part of modern software design. The technologies that enable cloud-native automation are built into the cloud platform on which cloud-native software runs. MNOs can use standard cloud-native tooling and approaches to build zero-touch automation that supports autonomous self-healing, upgrades, changes and performance optimizations of vRAN functions at scale. Such automation can massively reduce the operational complexity of a cloud-based RAN, if MNOs adopt an 'automation-first' mindset and bring discipline and focus to the task of creating it. If they do, MNOs should find opex reduction to be more achievable in a virtualized RAN context than it has been in other network domains in the past.

MNOs are looking for further sources of return on investment (ROI) beyond capex/opex savings for RAN virtualization. They are particularly interested in the service innovation that they expect virtualization to drive, led by technologies such as the near-real-time RAN Intelligent Controller (RIC), for example. The cloud-native architecture and open APIs that support the virtualization and disaggregation of the RAN should make it easier for MNOs to bring in and swap components, including cloud-native network function and platform components, software feature updates and security patches. This will allow them to take rapid advantage of vendor innovations to meet new customer service needs. MNOs expect a virtualized and disaggregated RAN, based on flexible, elastic and ubiquitous cloud infrastructure, to make them more agile in terms of service offers in future. In a digital economy, agility translates into competitive advantage.

It is noteworthy that automation is a prerequisite for achieving the top-5 drivers shown in Figure 2.2. This includes enabling new enterprise revenue streams (such as private 5G networks), where MNOs will need automation in the virtualized RAN to make them cost-competitive at new price points. Automation will also underpin customer self-service management of network slices. In NFV, it was possible to virtualize a network domain without automating it: in a cloud-native environment, full zero-touch automation and virtualization go hand-in-hand. The ability to build cloud-based automation is a beneficial side effect of taking an open, disaggregated approach to the RAN.

Figure 2.2: Commercial drivers for deploying vRAN before 2026⁴



2.3 Risks of deploying a disaggregated RAN

The following three types of risk are possible when virtualizing and disaggregating as large and complex a network as the RAN.

- Integration risk.** Disaggregation requires someone – often not the MNO – to re-aggregate and integrate components and to ensure that they all work together in a performant way, especially if each component is provided by a different vendor. Although open RAN initiatives mandate open interfaces between vCUs, vDUs and vRUs, these initiatives do not call for further interfaces to be standardised, for example, between the network functions and the underlying cloud infrastructure that they will run on, or with vendor-specific management tools. Such further integrations will require additional effort. If each vRAN component vendor has automated its component in an idiosyncratic way, integration time and effort will be needed to stitch automations together and remove any potential security vulnerabilities that have arisen from the use of different tools and approaches. MNOs will need to mitigate the risks associated with proprietary vendor interfaces and the use of proprietary tools and data management structures, which can make integration expensive and slow.
- Risk of vendor/platform lock-in.** MNOs should understand the granularity at which their prime contractors are disaggregating vendor functions in order to guard against new areas of lock-in. Lock-in might arise, for example, from choosing a single cloud provider to be the aggregator of an end-to-end RAN cloud platform and allowing that cloud provider to bring with it a limited set of vRAN network function vendors. A systems integrator or hardware vendor selected to build an end-to-end virtualized vRAN from multi-vendor components may nevertheless inhibit an MNO’s choice in a similar way to a single-vendor vRAN provider if it only works with a single pre-integrated vendor at each layer of the disaggregated, vRAN stack.

⁴ This data is taken from Analysys Mason’s survey of 74 MNOs concerning their vRAN and edge cloud architecture roadmaps, conducted in 3Q 2020.

- **Risks associated with management at scale.** Disaggregating the RAN takes place on a much larger scale than MNOs have experienced in any other network domain to date. A RAN is a highly distributed system, and disaggregated components need to be operated across many thousands of edge locations, rather than in a handful of data centers, as has been the case with other virtualized network functions, such as the mobile packet core or IMS. MNOs will also need their vRAN cloud infrastructure to operate in very different conditions than those that exist in central data centers, including supporting the highly demanding processing requirements of the vRAN's real-time functions (such as beamforming) on heterogeneous hardware with multiple accelerators. MNOs face a combinatorial management risk based on a number of sites, infrastructure criticality and a number of (disaggregated and re-composable) components in a vRAN.

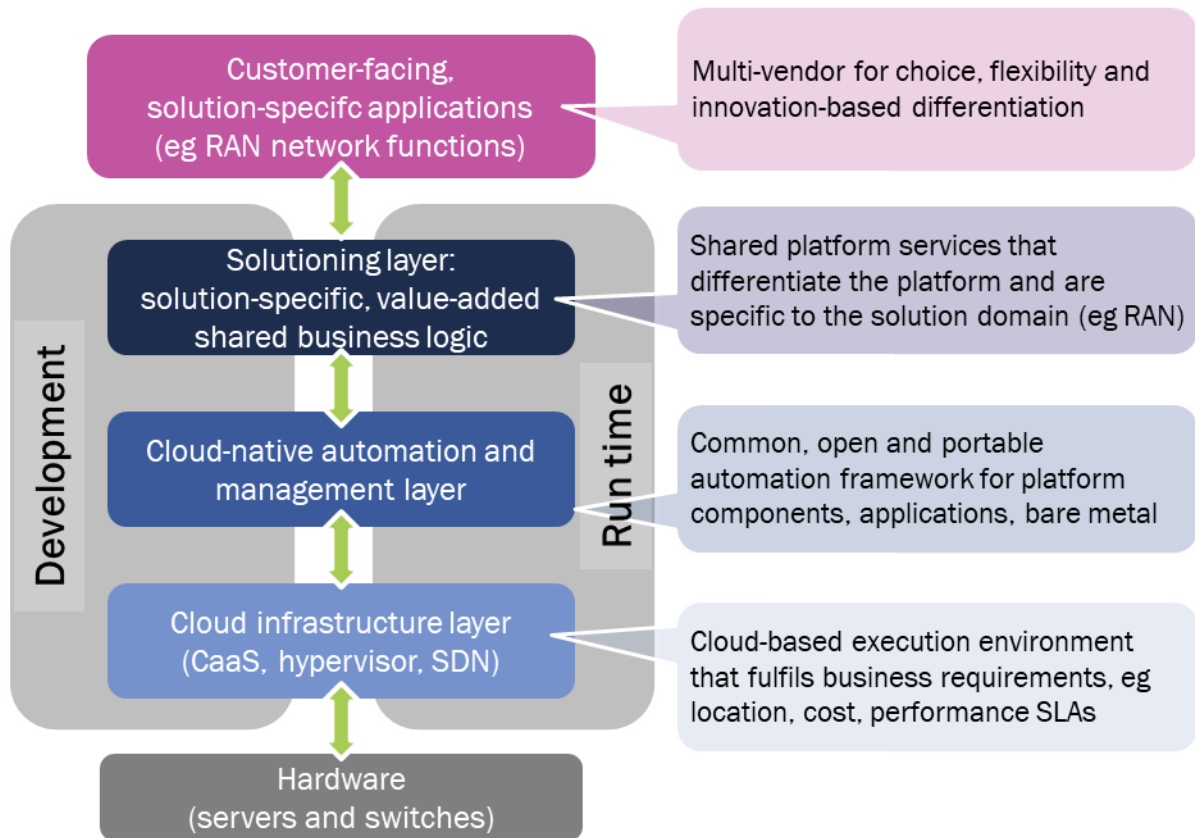
3. A platform architecture for structuring an open, virtualized RAN mitigates disaggregation risks

3.1 What is a platform?

Architecting a vRAN using a cloud-native platform approach addresses the risks of disaggregation and enables MNOs and systems integrators to develop management and automation in a consistent and reusable way across vRAN components.

Analysys Mason defines a platform as a structure that allows multiple applications participating in a 'solution' to be delivered within a common and shared technical and business framework. Platforms provide applications with access to shared assets (including common automation assets), enabling the highly efficient development, deployment and lifecycle operations of software-based solutions. Figure 3.1 shows the three layers of functionality in a cloud-based platform.

Figure 3.1: Three layers of platform functionality



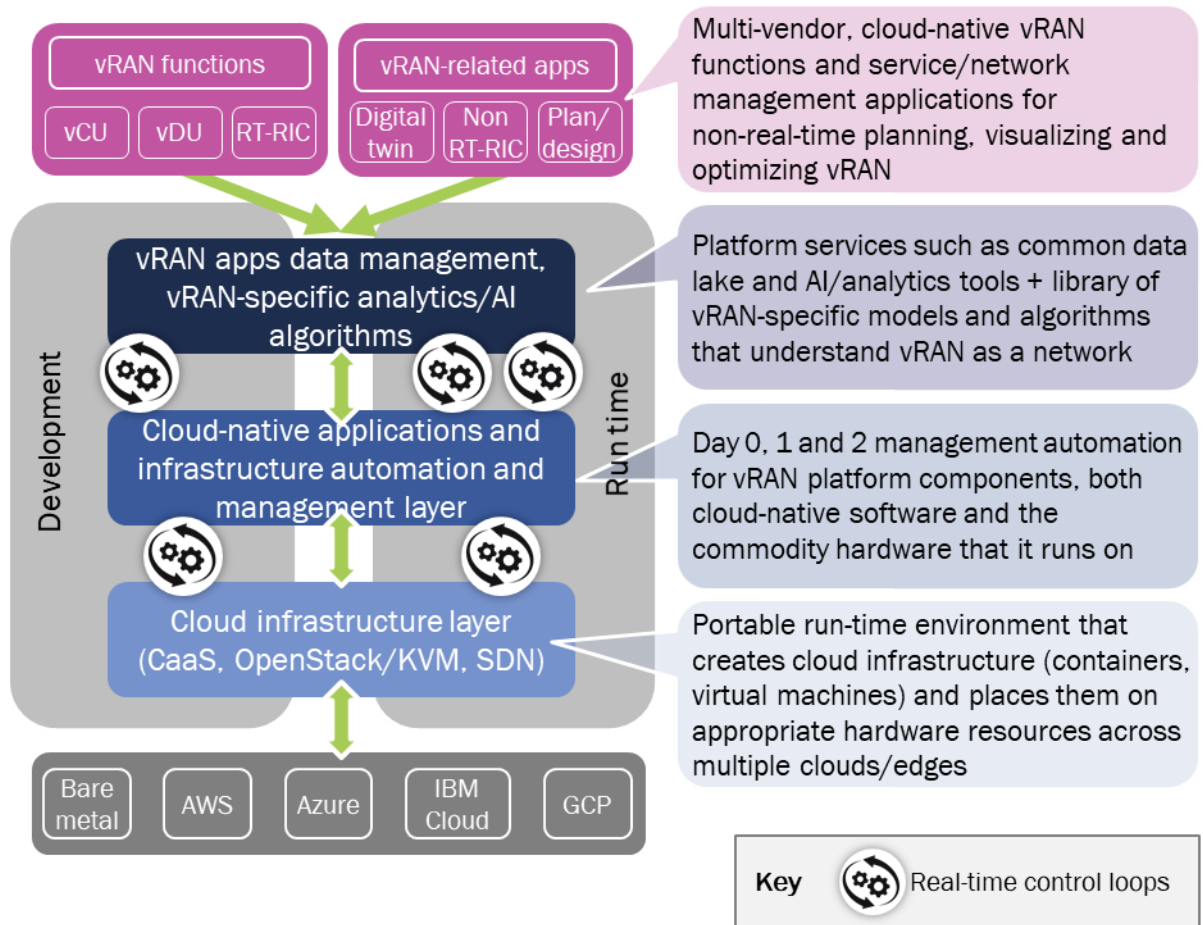
Source: Analysys Mason

Cloud-native vRAN solutions are platform-based: they consist of multiple network functions (the telecoms equivalent of IT applications), which run on top of a shared, technical framework (platform). Figure 3.1 illustrates the layers and components in a vRAN realized as an open, cloud-native solution.

In a cloud-native vRAN context, the bottom two layers of the platform use industry-standard, open-source components. The cloud infrastructure layer, for example, is based on Kubernetes and/or OpenStack, while the cloud-native automation and management framework contains Kubernetes-based tools developed within the Cloud-Native Computing Foundation. These include observability tools for the monitoring, logging and auditing of cloud-native software components in a vRAN solution (from its infrastructure components to its cloud-native network functions (CNFs)); and the Kubernetes Operator Framework, a toolkit for building operational automation (Day 1 and Day 2 configuration and management) into cloud-native software, such as CNFs.

The top layer of the platform contains the data structures and analytics/AI tools shared by all the CNFs that plug into the platform, regardless of vendor. Leading-edge MNOs that are implementing or have implemented a vRAN, point to the importance of having a data layer that is shared by all components in the vRAN solution because this provides a single source of truth for vRAN operation and management. **It is a striking feature of a cloud-native vRAN that every software component in it, whether in the platform or application layers, can be automated and managed in the same way using the same platform-based tools.**

Figure 3.2: A vRAN architected as a platform



Source: Analysys Mason

3.2 Platform-based automation mitigates disaggregation risks

MNOs will considerably reduce the risks associated with disaggregation if they select:

- vRAN functions that are built as software-only components and designed to run on such a platform architecture
- a cloud-native platform built from fungible, open and open-source components that can run across multiple cloud and bare metal infrastructures and with a common set of tools for building and executing zero-touch automation.

Integration risk is addressed by using platform-based common tooling, shared data and interfaces and an approach to building automation that is consistent across platform components and CNFs. **Lock-in risk** is reduced by architecting a vRAN solution as a set of well-bounded layers that can be flexibly implemented by using different products from multiple vendors. **The risks of management at scale** are mitigated through zero-touch automation, which carries out the same processes accurately and repeatedly across a highly distributed network. Because the automations are built in software, they can be modified and upgraded in software and applied remotely, which avoids the need for in-person visits to thousands of sites to make changes. Analytics and AI algorithms can be used to guide the automation, predicting outcomes and enabling proactive responses to emerging risks, wherever they appear. AI algorithms can be trained to understand the impacts of component

failure at different levels of the cloud platform on overall network performance, so that the right automated remedy can be applied when a specific failure event occurs. They will build knowledge of the optimal state of the network to ensure that different network components are contributing to that state, alerting orchestration systems to bring these components back into line if they diverge from it. The AI, analytics and automation tools underpinning a virtualized RAN are common, not only to its network functions, but to other applications as well, including business applications. This means that it will be easier to integrate RAN functionality and business processes in ways that have been difficult, if not impossible, to achieve in an automated way in the past. Operators will be able to use business intelligence to drive automated decisions at the network level, resulting in a virtualized RAN that is optimized for the business.

4. The benefits of cloud-native portability and zero-touch automation for a virtualized RAN

4.1 Designing automation for the virtualized RAN

Zero-touch automation needs to span three phases of a vRAN lifecycle: Day 0 deployment, Day 1 preparation for operation and Day 2 maintenance for the period in which the vRAN is live in production. We outline these phases in more detail below.

- Once a technician has plugged in the bare metal COTS hardware that will support the vRAN platform at a specific site, **Day 0 automation** controls the lighting-up of the hardware and the remote installation of the predefined, site-specific version of the vRAN platform software that needs to run there. This software consists of both the cloud infrastructure runtime (OpenStack and/or Kubernetes) and appropriate automation and management layer tools that will support the CNFs (such as security, monitoring and service mesh components). CNFs are pre-configured for a particular site and are also installed remotely.
- The entire, end-to-end Day 0 provisioning of a vRAN platform in a specific site should be governed by a **GitOps approach**. This uses a Git repository, in effect, a real-time inventory, which, in the case of a vRAN, holds all the site-specific definitions of the CNF/infrastructure software installed at each site. These site definitions, or models, will have been predefined in the Git repository, which is then used as the single source of truth driving each automated phased of vRAN lifecycle management. In the Day 0 phase, the appropriate site model in the Git repository is used to drive the automated installation of software ('infrastructure as code') in vRAN sites through a continuous integration/continuous delivery (CI/CD) workflow, or 'pipeline'. Automated workflows/pipelines are described using 'playbooks,' which can be enacted again and again without the risk of human error.
- The next stage of a GitOps-driven workflow then takes over. **Day 1 automation** tests and verifies all the vRAN components in their new environment, ensuring that all is ready for operation.
- Once all tests and validations have been conducted, the vRAN site is ready to go into production. Once live, **Day 2 automation** involves two means of managing vRAN sites in a zero-touch manner. GitOps-driven CI/CD pipelines are used to make changes to any type of vRAN software (for example, to apply patches and upgrades to platform components and/or CNFs). Real-time control loops can also be created to further self-heal and optimize vRAN components, keeping them available and performant. These real-time control loops can be built using platform tooling, including the Observability and Operator frameworks from the

management and automation layer. As MNOs become more experienced with closed-loop automation, they can augment their control loops with AI/analytics tooling from the solutioning layer of the platform to take preventative measures.

The following factors are important to the successful definition and implementation of zero-touch automation.

- All vendors that contribute disaggregated components to an overall vRAN solution will need to collaborate on its automation. Every vendor needs to commit to using the common tooling provided by the platform to build automation for their components, and must commit to a GitOps approach, with no manual or out-of-pipeline changes. This ensures that operators gain maximum value out of using a platform approach: more specifically, lower costs and complexity and easier integration of vendors because every vendor complies with an operator-specified common set of tools, rather than brings their own set.
- The development of real-time control loops is an iterative process, starting with feedback loops that are triggered manually in reaction to events and then progressing to closed loops that are triggered by scheduled events and expected traffic patterns. Once this level of autonomous automation has been mastered, an MNO will be ready to implement AI-driven control loops that can predict anomalous events and make appropriate changes to the vRAN to counter them.
- Zero-touch automation should be developed in a modular way and at a suitable level of abstraction to enable its reuse in other network domains. A vRAN pushes the boundaries of zero-touch automation: if it can be developed for such a demanding environment; it can also be applied in other areas of the network. The GitOps models will be different for vCPE and SD WAN, for example, but similar tools and pipelines should be used to automate their lifecycle management to achieve higher efficiency and reduce risk.

4.2 Benefits of automating a virtualized RAN

Virtualization of the RAN is in its infancy and very few operators have deployed one. Rakuten is one example of an operator that has deployed vRAN. It has been highly vocal about the benefits that it is seeing from automation, including faster and significantly less-expensive roll-out of its vRAN sites (which numbered 22 500 in May 2021) and a massive change in the number of software releases that it can absorb, from two releases a year to several a day.

Rakuten is, of course, a greenfield operator that has hired mostly ‘GenZ’ technical architects and software developers from across the world to build its vRAN automation in a cloud-native way using the GitOps-driven best practices described in this paper. However, established operators have had similar achievements in other network domains, paving the way for success in vRAN. An example is Deutsche Telekom, which set ambitious zero-touch automation targets for its platform-based Next Generation IMS project,⁵ in order to reach its target timescale of 3 months to bring features from specification to roll-out (compared to the 1–2 years, on average, that it takes in the traditional IMS environment; 2 days from the time a new release comes in from a vendor to roll it out to all live sites; 1 day to implement any patch; and zero nightshifts for maintenance). DT is now extending its ‘3-2-1-0’ vision to other areas of the network.

⁵ See Analysys Mason’s *A move to cloud changes the game for Deutsche Telekom’s next-generation IMS*. Available at: <https://www.analysismason.com/research/content/perspectives/dt-cloud-nims-rma16/>.

4.3 Selecting a platform for zero-touch automation

The vRAN platform provider will play a critical role in zero-touch automation because it will be the main provider of the tooling on which such automation will be built. MNOs should evaluate a platform provider based on the following criteria.

- Does it provide industry-leading, open-source DevOps tooling and cloud-native automation capabilities and is it actively involved in organizations that influence tools development, such as the Cloud Native Computing Foundation (CNCF)?
- Does it have a strong track record as a platform provider in the telecoms domain and has it worked with MNOs with advanced vRAN deployments and trials?
- How far does it support disaggregation? Can it provide a management and automation layer that is cloud infrastructure-agnostic, providing MNOs with choice at the infrastructure level?
- What is the extent of its vRAN ecosystem? Is it working with key partners in the ecosystem to co-develop automation playbooks and pipelines with them?
- What experience does it have of building zero-touch automation that addresses the timing, latency and other stringent vRAN-specific properties?
- Does it have platform-based AI and analytics capabilities that MNOs and vendors can use to increase the intelligence of their automations?

5. Conclusion

Virtualizing the RAN promises to bring multiple benefits, including the opportunities for service innovation and cost reduction. Automation is key to delivering these benefits and is enabled by the cloud-native platform technologies that will underpin a virtualized RAN. Operators should plan to leverage open and open-source tooling from the Kubernetes ecosystem from the start to build the zero-touch automation that they will need to reduce operational complexity. If they harness the automation capabilities built into the cloud-native platform on which all the components of a vRAN will run, they can reduce the risks associated with integration, vendor lock-in and management at scale. They can also gain significant agility, opex and security advantages as a result of using a common platform to automate and manage every software component in a consistent way, across cloud infrastructure, cloud-native network functions and even business applications.

Operators that create a collaborative environment around their common platform tooling and approach will find it easier to build the end-to-end automation that they need for zero-touch operation of the virtualized RAN. The different types of vendor that contribute functions to the vRAN (including cloud infrastructure, network function and service orchestration vendors) need to create automations for their specific software components that are compatible with those of other vendors. Operators will also gain the greatest benefit from vRAN automation if they standardize tools and lifecycle management approaches that can then be applied to other cloud-native network domains as these are deployed and integrated with the business environment. A platform that is open and based on widely used, industry-standard, open-source components (including AI and analytics tooling) will attract an ecosystem of vendors across network domains, boosting operators' ability to achieve their zero-touch automation ambitions.

6. About the author



Caroline Chappell (Research Director) heads Analysys Mason's Cloud research practice. Her research focuses on service provider adoption of cloud to deliver business services, support digital transformation and re-architect fixed and mobile networks for the 5G era. She is a leading exponent of the edge computing market and its impact on service provider network deployments and new revenue opportunities. She monitors public cloud provider strategies for the telecoms industry and investigates how key cloud platform services can enhance service provider value. Caroline is a leading authority on the application of cloud-native technologies to the network and helps telecoms customers to devise strategies that exploit the powerful capabilities of cloud while mitigating its disruptive effects.

This perspective was commissioned by Red Hat. Analysys Mason does not endorse any of the vendor's products or services.

Analysys Mason Limited. Registered in England and Wales with company number 05177472. Registered office: North West Wing Bush House, Aldwych, London, England, WC2B 4PJ.

We have used reasonable care and skill to prepare this publication and are not responsible for any errors or omissions, or for the results obtained from the use of this publication. The opinions expressed are those of the authors only. All information is provided "as is", with no guarantee of completeness or accuracy, and without warranty of any kind, express or implied, including, but not limited to warranties of performance, merchantability and fitness for a particular purpose. In no event will we be liable to you or any third party for any decision made or action taken in reliance on the information, including but not limited to investment decisions, or for any loss (including consequential, special or similar losses), even if advised of the possibility of such losses.

We reserve the rights to all intellectual property in this publication. This publication, or any part of it, may not be reproduced, redistributed or republished without our prior written consent, nor may any reference be made to Analysys Mason in a regulatory statement or prospectus on the basis of this publication without our prior written consent.

© Analysys Mason Limited and/or its group companies 2021.