




EVOLVING RAN ARCHITECTURES: EARLY LESSONS ON EMBEDDING FLEXIBILITY AND OPENNESS

The RAN is evolving towards greater openness and disaggregation, but the path is hybrid and incremental. This report explores the challenges operators face, early lessons from deployments, and why embedding openness and flexibility early is critical to long-term competitiveness.



Executive summary

The radio access network (RAN) is entering a new stage of evolution, as the industry moves beyond vision statements around open and disaggregated RAN into real-world deployments. Operators globally have aligned on the direction of travel: the future of RAN will be defined by greater openness and disaggregation. These shifts are no longer just aspirations; they are recognised as essential to reduce vendor dependency, accelerate innovation and bring cloud economics into the RAN domain.

Operators, vendors and ecosystem initiatives are now driving momentum through trials and early deployments, demonstrating tangible progress and creating a more pragmatic basis for adoption. Two themes stand out:

1. **Early challenges are being addressed, building confidence in adoption.** Advances in hardware acceleration are closing the performance gap with legacy systems, while joint testing and ecosystem collaboration are reducing the risk of interoperability issues in early stages. While these developments do not eliminate the complexity of achieving scale, they do demonstrate that technical and operational challenges can be overcome, giving operators greater confidence to move beyond pilots.
2. **The path is more hybrid and incremental than originally anticipated.** Many deployments today are blending some elements of the fully open, disaggregated, multi-vendor O-RAN vision, by opting for partial disaggregation – often through vendor-integrated virtualised RAN and mixes of public and private cloud environments. This approach enables operators to capture immediate benefits, such as reduced reliance on bespoke hardware and faster time to market, while retaining the accountability of a lead vendor to manage near-term integration risks.

However, this hybrid path underscores a critical tension. Prevalidated stacks and lead vendor integration can simplify deployments and may be a natural and necessary step to ensure end-to-end performance and accountability in early phases. However, this may entrench new forms of dependency over time. If telcos' operating models, skills base and lifecycle processes become too closely tied to a single vendor's integration blueprint, they risk recreating the very constraints that openness was meant to overcome.

This report explores the current state of RAN evolution and the growing momentum across the ecosystem. What is clear is that operators must **embed openness and flexibility at every stage** to maintain control over innovation cycles, supplier diversity and long-term network evolution. Our recommendations for operators are:

- **Embrace open, disaggregated infrastructure and horizontal orchestration to avoid new lock-in**
- **Prioritise open interfaces to preserve flexibility and optionality.**
- **Invest in hardware acceleration to deliver performance parity in the short term**
- **Architect the network for long-term modular upgrade and evolution.**

Table of contents

Executive summary 2

The RAN is moving towards open, disaggregated architectures 5

Early deployments are beginning to address initial challenges 7

 Concerns over technical, operational and commercial issues remain 7

 Real-world deployments are beginning to address key challenges 8

Flexibility and openness are critical to avoid lock-in and enable future innovation..... 10

 Two elements are emerging as central to avoiding this outcome: cloud flexibility and open
 interfaces 10

Key recommendations for operators 13

Table of figures

Figure 1: The hybrid evolution towards open, disaggregated architectures 5

Figure 2: True multi-vendor interoperability requires both vertical and horizontal openness 11

The RAN is moving towards open, disaggregated architectures

The future of the RAN is defined by greater openness and disaggregation. Operators globally are committed to this principle when considering future networks, as it will reduce dependence on single vendors and create room for faster innovation. Industry momentum is now firmly behind this vision, supported by global initiatives such as the O-RAN Alliance and growing ecosystems of vendors, integrators and hyperscalers.

Over the past few years, significant progress has been made in moving towards this vision. The industry has advanced beyond the initial concept stage, with growing vendor ecosystems and learnings from early trials and deployments. The challenge now lies less in 'if' networks will adopt greater openness and disaggregation, and more in 'how' and 'when' operators can make the transition at scale while retaining enough flexibility to evolve as architectures and technologies mature.

What has become clear is that adoption will be more hybrid and incremental than originally anticipated. Many deployments today are of partially disaggregated cloud RAN rather than fully open architectures: they combine elements of the full O-RAN vision adopted step by step, often through vendor-integrated virtualised RAN and mixes of public and private cloud environments. The fully open, disaggregated, multi-vendor blueprint— spanning radio unit (RU), distributed unit (DU), centralised unit (CU) and orchestration — remains aspirational for large-scale rollouts, given the complexity of achieving full interoperability. Not all operators will aim for this as their end state of RAN evolution; instead, many are prioritising partial openness such as Open Fronthaul, while continuing to rely on lead vendors to ensure integration and performance.

Figure 1: The hybrid evolution towards open, disaggregated architectures

	Legacy	Network evolution	
	Traditional RAN	Partially disaggregated cloud RAN	Fully open, disaggregated RAN
Virtualisation	✖✖ Hardware and software tightly coupled in proprietary stack	✓✓ Software-virtualised and able to run on COTS servers (e.g., CU/DU separation)	✓✓ Full virtualisation across RAN functions, with standardised virtual infrastructure
Cloud-native	✖✖ Not cloud-native, tied to physical appliances	✓✓ Container-based architectures (CNFs, CaaS, Kubernetes)	✓✓ Fully cloud-native, with microservices and CI/CD automation
Open interfaces	✖✖ Proprietary interfaces	✓ Partial openness (e.g., via open fronthaul or management interfaces)	✓✓ Standardised interfaces enabling interoperability
Disaggregation	✖✖ RAN components deployed together as integrated BBU stack	✓ Initial logical and functional splitting of RAN components	✓✓ RAN components (RU/DU/CU) can be procured and deployed in separate geographic locations
Multi-vendor interoperability	✖✖ End-to-end single-vendor only	✓ Partial multi-vendor compatibility, mainly at RU-DU boundary	✓✓ Multi-vendor mix and match across RU, DU, CU and orchestration

Source: STL Partners

Each approach offers tangible benefits. Moving from traditional RAN to partially disaggregated cloud RAN reduces reliance on bespoke hardware, introduces cloud economics – e.g., scalability, cost efficiency across commercial off-the-shelf (COTS) and cloud platform, accelerates operator autonomous network journeys and lays the foundation for more open architectures. At the same time, it allows operators to continue support from a single vendor for end-to-end integration, ensuring performance and accountability in the near term.

Deployments that extend further towards a fully, open disaggregated RAN provide greater flexibility and control, accelerate innovation (as they are no longer tied to single-vendor innovation cycles) and enable true multi-vendor interoperability.

However, as operators continue this hybrid path, the degree of openness and flexibility they build into infrastructure will determine how much control they retain and how easily they can diversify suppliers over time. Avoiding new forms of lock-in – whether to vendors, integrators or cloud providers – will be essential to ensure today's architectures do not constrain future evolution and innovation.

Looking further ahead, 6G is expected to follow this trajectory. While the precise model is not yet defined, it will almost certainly build on openness, disaggregation and programmability, while introducing new AI-native control functions and deeper automation across the RAN, core and edge. In practice, this may not replicate the open RAN model exactly (as it has been envisioned today by the O-RAN Alliance), but it will continue its trajectory towards more modular, software-driven and intelligent networks.

This report explores the current state of RAN evolution and the growing momentum across the ecosystem. Whilst there remain several challenges which operators must navigate, there is a shift in the industry, as early deployments are beginning to demonstrate how these issues can be addressed in practice. As operators evolve on their hybrid path, they must embed openness and flexibility early to ensure long-term competitiveness and control.

Early deployments are beginning to address initial challenges

Concerns over technical, operational and commercial issues remain

Whilst operators' commitment to more open, disaggregated RAN architectures is strong, progress has inevitably been shaped by the realities of operating large-scale networks. Rather than a discrete shift, most operators are taking an incremental approach – deploying partially disaggregated and fully open RAN components alongside traditional, purpose-built RAN.

Despite demonstrated progress, concern remains with the practical challenges in translating these new architectures into scaled, real-world deployments:

- **Technical challenges:** Delivering performance parity with legacy RAN continues to be a critical benchmark. For operators adopting more open, disaggregated architectures, this means ensuring that virtualised, cloud-native baseband functions can achieve equivalent throughput, latency and energy efficiency to traditional systems. Managing hybrid networks also increases complexity, as integration and orchestration must span diverse hardware, cloud platforms and software releases.
- **Operational challenges:** Moving away from purpose-built, pre-integrated RAN introduces new layers of complexity. Operators now need to procure and deploy COTS hardware at scale, select and manage the cloud platform and ensure integration across these layers – often with multiple new vendor relationships to manage. While some multi-vendor interoperability exists today, most partially disaggregated cloud RAN deployments allow operators to rely on a single vendor to reduce this complexity. The shift towards a fully open, disaggregated RAN takes one step further: enabling multi-vendor interoperability within the RAN cluster itself (e.g., radios and baseband components from different suppliers). This brings added complexity for integration, lifecycle management and accountability. At the same time, operators continue to highlight internal skills gaps: cloud-native practices such as CI/CD pipelines and Kubernetes-based orchestration remain unfamiliar to traditional network engineering teams, slowing the operational transition.
- **Commercial challenges:** The business case for open, disaggregated architectures is still evolving. This shift requires upfront investment in infrastructure, cloud platforms and integration services. Cost efficiencies depend on reaching scale – and some specialist vendors still lack the volume to match incumbent suppliers' pricing power.

These challenges are not new and have been central to the RAN evolution debate from the beginning. However, what has changed is the availability of tangible deployment experience, providing operators with greater clarity of how performance, integration and operational risks can be managed. This shift is creating a more pragmatic basis for adoption, accelerating progress and creating a clearer direction of travel.

Real-world deployments are beginning to address key challenges

While the journey towards fully open, disaggregated RAN architectures remains a work in progress, the past two years have seen a noticeable shift away from traditional RAN and towards the characteristics of 'open' through deployments of partially disaggregated cloud RAN and small-scale deployments of O-RAN. Challenges around performance, integration and operations are no longer only theoretical: they are being tested in the field, with vendors and operators generating the first wave of evidence on what works. These early deployments remain limited in scale, but they are important in demonstrating pathways to reduce integration complexity, achieve performance parity and support day-to-day operations in hybrid environments.

1. Pre-integrated solution components, blueprints and validation

One of the most pressing barriers to adopting more open architectures has been the integration burden on operators. To address this, vendors are delivering reference designs and blueprints that combine hardware, cloud platforms and software in prevalidated stacks. Developed through joint test labs, vendor alliances and hyperscaler partnerships, these approaches help operators accelerate commercial readiness while reducing the risks of multi-vendor interoperability. Some are going further by offering pre-integrated infrastructure components, enabling communication service providers (CSPs) to offload much of the testing and validation effort.

This allows operators to realise the benefits of more open, disaggregated architectures while continuing to rely on single-vendor accountability for performance and integration.

Industry examples

- In the UAE, du deployed its first commercial 5G standalone (SA) Cloud RAN site in Abu Dhabi using Nokia's anyRAN solution, running virtualised baseband functions on Dell PowerEdge XR8620 servers with Red Hat OpenShift. By relying on Nokia's pre-integrated reference designs, du was able to validate interoperability across distributed and central units, radios and orchestration layers, enabling the rollout of a hybrid RAN without extensive in-house testing.
- In Finland, Elisa took a similar approach with a commercial site in Espoo, using the same Nokia–Dell–Red Hat blueprint. The deployment was tested against benchmarks for capacity, latency and automation, demonstrating that prevalidated stacks can shorten time to market while maintaining service performance.

2. Performance parity through hardware acceleration

Performance parity with traditional RAN has long been the benchmark for credibility. Progress is now visible, with leading operators running live deployments of partially disaggregated cloud RAN, demonstrating significant improvements in performance KPIs. As part of this, Layer 1 (L1) acceleration is a key technical enabler. Industry efforts are converging around two different acceleration solutions: in-line and look-aside acceleration. These solutions allow partially disaggregated cloud RAN workloads to achieve comparable throughput, latency and energy efficiency to purpose-built equipment.

Industry examples

- Nokia, working with Marvell Technology, has developed an in-line L1 acceleration SmartNIC integrated into Dell servers. Trials with Elisa in Finland and e& in the UAE showed that cloud RAN workloads achieved comparable downlink and uplink rates, with latency in line with purpose-built RAN, while also improving energy efficiency per cell. Crucially, operators reported that scaling capacity no longer required linear increases in CPU resources, addressing concerns about total cost per cell.
- These results suggest that hardware acceleration can move performance concerns from a theoretical risk to a manageable engineering challenge. For operators, the implication is that open RAN can begin to meet the KPI standards required for scaled deployment, rather than being confined to low-traffic or greenfield environments.

3. Day-2 operational support in hybrid environments

Operational resilience is another critical concern, as CSPs experiment with hybrid networks combining purpose-built and partially disaggregated cloud RAN. A key challenge they face is skills to manage the cloud-native environment alongside legacy architectures. Vendors are responding by extending their orchestration and management platforms to cover day-0, day-1 and day-2 operations, including automation, upgrades and performance optimisation. These services are especially important for operators that do not yet have extensive in-house cloud-native expertise.

Industry examples

- Nokia has extended its MantaRay network management portfolio to cover both purpose-built and cloud RAN. In operator pilots, MantaRay delivered automated lifecycle management, upgrades and fault monitoring across CU, DU and RU, reducing the need for manual intervention.
- In parallel, a demonstration at Mobile World Congress Barcelona 2024 by Dell, Red Hat and Amdocs showcased AI-driven closed-loop automation running on prevalidated telecom infrastructure. The demo showed end-to-end observability and zero-touch lifecycle management of cloud RAN functions, highlighting how day-2 operations can be simplified even in multi-vendor environments.

Flexibility and openness are critical to avoid lock-in and enable future innovation

As operators transition towards fully open disaggregated RAN, immediate priorities often centre on simplifying integration. Prevalidated stacks and blueprints reduce deployment risk in the short term, and initial reliance on a lead vendor for integration can be a natural and often necessary step to ensure end-to-end performance and accountability in early phases. However, this may risk creating new forms of dependency over time, if an operator's operational mode, skills base and lifecycle management process become too tightly reliant on the lead vendor. For example:

- At the RAN function layer, most deployments today still rely on a lead vendor for CU/DU initial integration, which may limit future multi-vendor innovation inside the RAN cluster.
- At the orchestration and containers-as-a-service (CaaS) layers, operators may risk being tied to a single vendor's lifecycle management system or Kubernetes distribution, constraining their ability to adopt alternative cloud platforms.
- At the cloud infrastructure layer, overreliance on a single public cloud or private cloud platform risks repeating the vendor concentration of legacy networks.
- In practice, the risk is not just 'vendor lock-in' but also systems integrator (SI) lock-in, as many operators are turning to prime vendors to assemble and certify end-to-end stacks. While this reduces friction in the short term, it raises concerns about replicating ongoing dependency models from legacy RAN, if operators do not build the skills to manage the cloud-native operations.

Two elements are emerging as central to avoiding this outcome: cloud flexibility and open interfaces


A growing number of deployments are running workloads across multiple virtualisation layers and cloud platforms, demonstrating that multicloud hosting is becoming increasingly viable.¹ Cloud flexibility – the ability to run RAN workloads across both public and private clouds, and across different CaaS platforms (e.g., OpenShift, Wind River) – is increasingly critical for long-term control and future vendor flexibility.

This is also reinforced by the trend towards horizontal, cloud-native infrastructure: common platforms that can support multiple telco workloads (RAN, core, edge) on the same hardware and orchestration system. This approach reduces hardware interoperability complexity and allows operators to evolve networks in a more modular way. Instead of siloed stacks, operators can add or replace workloads incrementally, avoiding technical debt that might otherwise constrain their path towards 6G.

¹ Covered in our recent report: ['Telco cloud deployment tracker Q1 2025: The various flavours of open RAN'](#).

Yet, cloud flexibility on its own is insufficient if interfaces between RAN components and management layers remain closed. Vertical openness (through disaggregation of hardware and software) may enable cloud flexibility, but without horizontal openness (through interfaces between RAN functions and management layers), the risk of intra-RAN lock-in remains.

Figure 2: True multi-vendor interoperability requires both vertical and horizontal openness



Source: STL Partners

Whilst full multi-vendor disaggregation is not yet common, the growing support of open interfaces enables important steps towards long-term flexibility and control. For example:

- Open Fronthaul (O-RAN 7.2x) has enabled operators to begin sourcing radios independently of DU/CU software, fostering greater vendor diversity and reducing dependency on fully proprietary hardware-software stacks.
- The O2 interface, defined by the O-RAN Alliance, is gaining traction as a way to standardise interactions between RAN management systems and underlying cloud infrastructure (O-Clouds). This allows operators to avoid dependency on a single cloud vendor by enabling lifecycle management across different CaaS and infrastructure platforms.
- As the ecosystem matures, RAN intelligent controller (RIC) interfaces (E2, A1, O1) will allow operators to bring in third-party optimisation and AI-driven control apps. These capabilities would be much harder to integrate if orchestration remained under the exclusive control of a single lead vendor.

Despite the existence of open interfaces making multi-vendor interoperability more technically feasible, most deployments today still rely heavily on vendor-led integration. This is understandable given the immaturity of ecosystems and the pressure to deliver quickly. But over time, operators risk swapping one form of dependency (traditionally integrated RAN) for another (vendor- or SI-integrated

RAN stacks). By embedding infrastructure flexibility early, operators can adapt as 6G architectures emerge, avoiding technical debt that limits innovation and competitiveness. In practice, this shift won't be defined solely by O-RAN Alliance standards. Operators and vendors will selectively adopt the interfaces that deliver the greatest value – some aligned with O-RAN specifications, while others evolving towards greater openness and disaggregation in a different way.

Key recommendations for operators

- Embrace open, disaggregated infrastructure and horizontal orchestration to avoid new lock-in.** As operators shift away from pre-integrated systems, adopting disaggregated hardware and software supported by horizontal orchestration and automation layers will be essential to maintain long-term flexibility. This approach enables end-to-end lifecycle management, reduces capital expenditure (capex) and operating expenses (opex) through cloud economics, and prevents recreating the dependency models of legacy RAN.
- Prioritise open interfaces to preserve flexibility and optionality.** Even where operators rely on vendor-integrated deployments today, prioritising open networking and open interfaces – such as Open Fronthaul (7.2x) and O2 – will create options for introducing new vendors or components over time. Open networks not only support vendor diversity but also enable the integration of future functionality, including AI-native applications and RIC-based control, without the need for wholesale platform replacement. Partnering with vendors that actively push for fully open and disaggregated stacks (e.g., participate in the O-RAN Alliance and demonstrate real-world multi-vendor interoperability) will accelerate this flexibility.
- Invest in hardware acceleration to deliver performance parity in the short term.** Performance parity with purpose-built RAN is a prerequisite for scaling partially disaggregated cloud RAN or moving further towards fully open, disaggregated RAN in demanding environments such as urban massive MIMO. Hardware acceleration has now been proven to close this gap, allowing workloads to deliver comparable throughput, latency and energy efficiency to integrated systems. Operators should ensure their platforms are designed to incorporate in-line acceleration while remaining compatible with COTS hardware and cloud-native models, avoiding performance bottlenecks that could limit commercial viability.
- Architect the network for long-term modular upgrade and evolution.** Operators should design RAN platforms with flexibility built in, enabling the disaggregation of RU, DU and CU functions, so that components can be upgraded or swapped independently as supplier ecosystems evolve. To avoid cloud infrastructure lock-in, cloud RAN software should be deployable across multiple public and private cloud environments, as well as CaaS platforms. Looking ahead, networks should also be prepared to integrate RIC-based AI-native applications (xApps/rApps), as programmability and automation will become central to 6G-era architectures.

PARTNERS



Research



Consulting



Events