The Pragmatic Openness of Open RAN: Navigating the Path from Vendor Lock-in to Interoperability Islands

Nima Afraz, Mohammad Shojafar, *Senior Member, IEEE,* Johann M. Marquez-Barja, *Senior Member, IEEE,* and Hamed Ahmadi, *Senior Member, IEEE*

***Abstract*—Open Radio Access Network (Open RAN) promises to dismantle traditional vendor lock-in through disaggregation and standardization. This article analyzes its current trajectory, finding that the ecosystem is not evolving toward universal interoperability but is instead consolidating into "islands of interoperability" or pre-validated multi-vendor clusters that risk becoming new forms of lock-in. Through a review of global deployments and industry collaboration patterns, we examine the technical and market forces that may contribute to this fragmentation. We argue that while these islands represent a pragmatic compromise that increases operator choice, they also pose a significant long-term risk to the open ecosystem by stifling innovation and competition. Ultimately, we conclude that avoiding permanent fragmentation is not a technical challenge alone but a strategic imperative. It requires a concerted shift from writing standards to enforcing rigorous, transparent certification and fostering a market that rewards true interoperability over closed partnerships. The future of Open RAN depends on building bridges between these islands to preserve its foundational promise.**

Link to graphical and video abstracts, and to code: https://latamt.ieeer9.org/index.php/transactions/article/view/10297

***Index Terms*—Open RAN, vendor lock-in, disaggregation, interoperability, certification, multi-vendor, network slicing, AI/ML.**

I. Introduction

RADIO ACCESS NETWORKS (RANs) have historically been proprietary and vertically integrated, where one vendor delivers the radio units, baseband hardware, and control software as a closed solution. This model can create significant vendor dependence, limiting operators' flexibility and reducing supplier diversity. Open RAN was introduced as a paradigm shift aimed at countering this lock-in by disaggregating the RAN architecture and defining open, standards-based interfaces between its components [1]. Recent studies have highlighted that Open RAN not only reduces capital and operational expenditures but also improves network performance metrics such as latency and throughput. Aljanabi et al. suggest through simulation and case studies with operators like Rakuten, Vodafone, and Telefónica that Open RAN architectures can achieve up to 25% higher throughput and 20% lower latency compared to traditional RAN systems, confirming its potential to decentralize telecom networks and enhance sustainability [2] In Open RAN, the baseband processing units (Centralized Unit (CU), Distributed Unit (DU)) and the Radio Unit (RU) and the x/rApps can be supplied by different vendors, provided that they adhere to common interface specifications. This disaggregation enables operators to mix and match components, thus enabling competition and innovation. Furthermore, Open RAN embraces virtualization and cloudification. Virtualizing RAN functions (vRAN) allows for deployment on general-purpose hardware, reducing costs and increasing scalability. Additionally, the integration of a RAN Intelligent Controller (RIC) introduces native support for Artificial Intelligence (AI) and Machine Learning (ML), enabling dynamic, data-driven network optimization [1]. Governments and industry alliances have also embraced Open RAN for supply chain diversity – for example, the UK set a goal that 35% of the country's mobile traffic should ride on open/interoperable RAN by 2030. Major carriers in Europe (Deutsche Telekom, Orange, Telefónica, Vodafone, etc.) signed an MoU in 2021 committing to make Open RAN the "technology of choice" for future networks . Amid this enthusiasm, how "open" is Open RAN in practice today? After several years of trials and initial deployments, a gap has become evident between the ideal of fully plug-and-play multi-vendor networks and the reality of what's achievable with current technology and market conditions. This article explores that question, examining the architectural vision, real-world implementations, market implications, case studies from around the globe, and what experts foresee for Open RAN's future. Beyond market dynamics and policy implications, this paper contributes a technical analysis of why Open RAN interoperability remains partial in practice. Specifically, it (i) formalizes the concept of *islands of interoperability* as an architectural manifestation of bounded multi-vendor compatibility; (ii) applies a graph-theoretic co-occurrence analysis with

The associate editor coordinating the review of this manuscript and approving it for publication was Carolina Del-Valle-Soto (*Corresponding author: Nima Afraz*).

This work was supported in part by Taighde Éireann–Research Ireland under Grant 13/RC/2077_P2 through CONNECT, the Research Ireland Centre for Future Networks, and in part by the European Union's Horizon Europe research and innovation programme through the Marie Skłodowska-Curie SE project RE-ROUTE under Grant Agreement No. 101086343.

Nima Afraz is with the School of Computer Science, University College Dublin, Ireland (e-mail: nima.afraz@ucd.ie).

M. Shojafar is with the 5G/6GIC, Institute for Communication Systems, University of Surrey, United Kingdom (e-mail: m.shojafar@surrey.ac.uk).

J. M. Marquez-Barja is with the IDLab, University of Antwerp - imec, Belgium (e-mail: johann.marquez-barja@uantwerpen.be).

H. Ahmadi is with the Department of Electronic Engineering, University of York, United Kingdom (e-mail: hamed.ahmadi@york.ac.uk).

Louvain community detection to 31 deployment and ecosystem activities; (iii) identifies the principal technical drivers of fragmentation at the interface and RIC levels, including integration overhead, timing constraints, xApp conflicts, and certification limitations; and (iv) presents a structured 2025–2035 roadmap linking these challenges to concrete technical enablers and mitigation paths.

This paper is intended as a technical analysis and perspective article supported by empirical evidence, rather than as a controlled experimental study. Its contribution is to examine how Open RAN openness is realized in practice by combining architectural discussion, ecosystem evidence, and co-occurrence-based analysis of publicly reported Open RAN activities.

II. Islands of Interoperability

An *island of interoperability* is a cluster of technologies that interoperate across a limited set of vendors while remaining isolated from a fully open ecosystem. Standards may exist, but political and economic incentives prevent complete plug-and-play compatibility. Electric-vehicle charging illustrates this dynamic. Competing standards such as CCS, CHAdeMO, Tesla's NACS, and China's GB/T all deliver the same service but differ in connectors, protocols, and backend control. Tesla built its own network to lock in users and capture data revenues, European regulators promoted CCS to protect domestic automakers, Japan advanced CHAdeMO to maintain industrial influence, and China mandated GB/T to secure national control. Each actor seeks to safeguard revenue streams, technological sovereignty, or bargaining power, so drivers encounter fragmented networks that require adapters or roaming agreements despite technical possibilities for universal access.

Open RAN faces similar forces. Although O-RAN specifications define open interfaces, large vendors certify only selected partner combinations, while governments promote national implementations to reduce reliance on rivals. Control of data and proprietary algorithms in the RAN Intelligent Controller reinforces partial openness. As in EV charging, bridging these islands demands not only technical compliance but also the alignment of commercial and geopolitical interests.

III. Open RAN Architecture Vision vs. Legacy RAN

Fig. 1 contrasts three network architectures: a *traditional* legacy RAN, a *fully interoperable* Open RAN, and the *practical islands of interoperability* that characterize many current deployments. In a **legacy RAN** (left panel), the radio unit (RRU), baseband unit (BBU), and operations support system (OSS/EMS) form a vertically integrated stack from a single vendor. Internal interfaces are proprietary and tightly optimized for performance. This "one-vendor" model simplifies procurement and accountability, operators have "one throat to choke" when a fault occurs, but it locks operators into a single supplier and limits their ability to introduce innovations from other vendors. With only a few companies able to provide complete RAN solutions, and with geopolitical restrictions further shrinking the pool of vendors, operators began to seek a more open architecture. A **fully interoperable Open RAN** (center panel) aims to disaggregate *all* major RAN functions, not only the radio (O-RU), distributed unit (O-DU), and central unit (O-CU), but also the RAN Intelligent Controller (RIC) and its applications (xApps in the near real-time RIC and rApps in the non-real-time RIC), so that each can be sourced from different vendors and combined through standardized interfaces. The O-RAN Alliance defines these open interfaces, including fronthaul (O-RU to O-DU), midhaul (O-DU to O-CU), E2 (near-RT RIC to RAN nodes), A1 (non-RT to near-RT RIC), O1 (SMO to network functions), and O2 (SMO to cloud infrastructure). In this "Lego-style" architecture, an operator might deploy radios from Vendor A, baseband units from Vendor B, a near-RT RIC from Vendor C, and a portfolio of xApps and rApps from independent third parties, all orchestrated by a cloud-native Service Management and Orchestration (SMO) layer. Extending multi-vendor openness to the application layer enables specialized algorithms for traffic steering, anomaly detection, or energy optimization to be procured competitively and updated independently of the underlying hardware. The right panel illustrates the **current reality**: *islands of interoperability*. Early Open RAN deployments often cluster around vendor-specific subsystems, each "island" contains a near-RT RIC, non-RT RIC, CU/DU, and RU that interoperate internally but rely on limited, pre-tested cross-vendor combinations. While these implementations conform to O-RAN specifications, they stop short of the fully mix-and-match ecosystem envisioned by the standard. Operators can combine vendors at the fronthaul or midhaul, but application portability across different RIC platforms (for xApps and rApps) remains a work in progress, and performance parity with single-vendor systems requires significant integration effort. Early proof-of-concept trials and laboratory demonstrations have validated that multi-vendor interoperability is technically possible. Industry groups have hosted plugfests, such as the Telecom Infra Project (TIP) labs and the i14y Lab in Berlin, to test equipment from different suppliers. Major vendors including Nokia and Ericsson have introduced O-RAN-compliant products; for example, Nokia's 5G baseband software supports the O-RAN fronthaul specification and has been integrated with third-party radio units. Nevertheless, realizing a fully disaggregated RAN that matches the performance and feature set of a single-vendor system remains challenging. The RAN is a real-time system requiring tight coordination for scheduling, beamforming, and handovers. Introducing multiple vendors and open interfaces increases integration complexity and adds potential overhead unless carefully optimized. In practice, many initial Open RAN deployments consist of vendor-specific subsystems, such as one vendor's O-DU paired with its own O-RU, that are O-RAN compliant but not yet fully mix-and-match. While the vision of complete openness is compelling, live network implementation remains a complex engineering task, and the next sections examine how "open" current Open RAN deployments truly are and what technical and operational limitations have emerged.

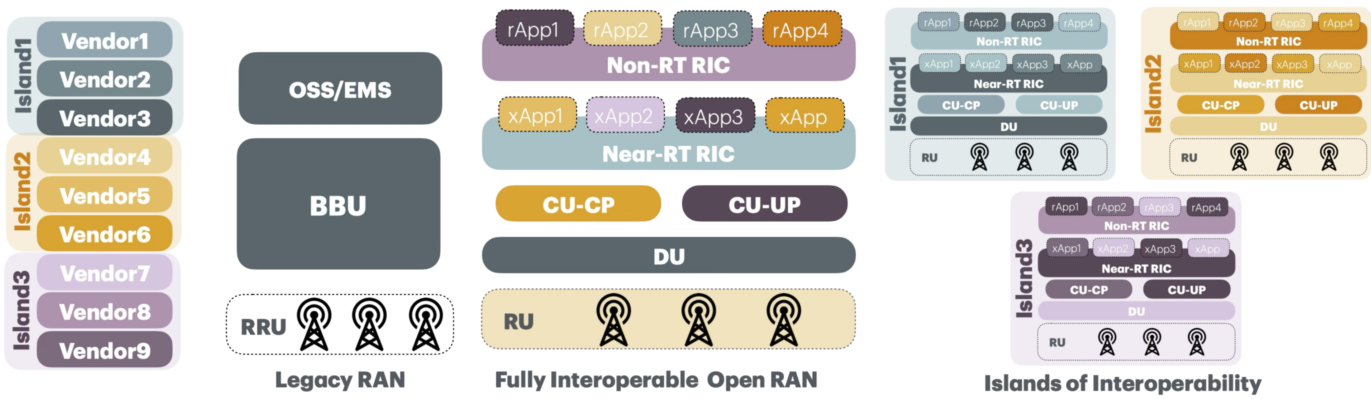

Fig. 1. Traditional RAN, fully interoperable Open RAN, and the practical islands of interoperability observed in today's deployments.

A. Defining and Testing Interoperability

The IEEE defines *interoperability* as "the ability of two or more systems or components to exchange information and to use the information that has been exchanged" [3]. In the context of Open RAN, this basic notion is extended to the disaggregated radio access network where multi–vendor hardware and software must not only exchange control and user plane data, but also maintain deterministic timing, performance, and security.

The O-RAN ALLIANCE treats interoperability as a core requirement for realizing openness at scale. Commercial success depends on "scale and interoperability," and that openness, intelligence, and virtualization remain the three fundamental pillars of O-RAN [4]. To operationalize this goal, the alliance specifies open interfaces (e.g., F1, E2, A1, O1/O2, Open Fronthaul) and develops detailed interoperability test specifications. Empirical evidence from early real-world deployments reinforces the challenges inherent in realizing full openness within multi-vendor O-RAN ecosystems. Krasniqi et al. [5] provide first-hand observations from private 5G testbeds, such as SONIC Labs and Vodafone's validation facilities, showing that even with well-defined F1, E2, and A1 interfaces, integration remains hindered by inconsistent vendor implementations, manual configuration overhead, and limited support for automated CI/CD-based testing. Vendor interoperability within Open RAN introduces challenges beyond those in traditional, vertically integrated systems due to the multiplicity of interfaces and stakeholders involved [6]. Nonetheless, the O-RAN Alliance has embedded several mechanisms that ease integration, such as standardized E2-based discovery and capability-exchange functions, detailed conformance testing procedures, and a distributed network of more than fifteen Open Testing and Integration Centers (OTICs) worldwide. Verification is supported by a global network of OTICs, which provide common laboratories and procedures for functional verification, integration, badging, and certification. Interoperability testing defined by O-RAN Working Group 5 ensures that multivendor components operate seamlessly across open interfaces by validating key interactions such as CU–DU communication on the F1 interface and RU–DU connectivity on the Open Fronthaul link. Through these standardized processes, the alliance aims to reduce integration risk and enable operators to deploy disaggregated RAN solutions with predictable performance. Table I shows that Open RAN activity is not uniform across deployments, actors, or technical objectives. Some initiatives focus primarily on radio access disaggregation and O-RU/O-DU integration, while others emphasize RIC-based programmability, cloud-native deployment, private networks, or large-scale operator trials. This variation indicates that "Open RAN deployment" is not a single technical category, but a spectrum of integration scenarios with different interoperability requirements. The comparison also supports the notion of emerging interoperability islands. Several pilots and deployments involve recurring combinations of operators, vendors, integrators, and cloud providers, suggesting that practical interoperability often develops first within pre-tested collaboration groups. As a result, openness at the architectural or standards level does not automatically imply universal plug-and-play substitutability across all vendors. Instead, Table I illustrates how interoperability remains shaped by deployment context, vendor roles, integration maturity, and the specific interfaces being validated.

B. The Promise of Multi-Vendor Openness

Open RAN envisions a disaggregated radio access network in which *all key components*, the radio unit (O-RU), distributed unit (O-DU), central unit (O-CU), RAN Intelligent Controller (RIC), the applications running on it (xApps in the near-real-time RIC and rApps in the non-real-time RIC), and the Service Management and Orchestration (SMO) layer, can be sourced from different vendors and interconnected through standardized interfaces such as fronthaul (O-RU to O-DU), midhaul (O-DU to O-CU), E2 (near-RT RIC to RAN nodes), A1 (non-RT to near-RT RIC), O1 (SMO to network functions), and O2 (SMO to cloud infrastructure) [1]. This "Lego-brick" model allows, in principle, a radio from Vendor X, a baseband from Vendor Y, a near-RT RIC from Vendor Z, and third-party xApps or rApps to be combined and orchestrated on generic cloud hardware, enabling competition and lowering costs. Governments and industry coalitions, including the Open RAN Policy Coalition (ORPC) and the Global Coalition on Telecommunications (GCOT), have embraced this vision to diversify supply chains and explicitly warn against the

TABLE I
SELECTED OPEN RAN FIELD DEPLOYMENTS WORLDWIDE

Operator / Region	Year	Key Vendors	Deployment Scope	Model	Notable Outcomes
Rakuten Mobile, Japan	2020	Altiostar (Rakuten Symphony), NEC, Fujitsu, Cisco	Nationwide greenfield 4G (Apr 2020); 5G added Sep 2020; >50,000 sites	**Full multi-vendor**	World's first nationwide cloud-native Open RAN; >97% population coverage
DISH Wireless, USA	2022	Mavenir, Altiostar/Rakuten Symphony, Fujitsu, MTI, Nokia, AWS, Samsung	Nationwide 5G greenfield on AWS cloud; >15,000 sites by June 2023	**Full multi-vendor**	70% U.S. population coverage achieved June 2023; first U.S. cloud-native O-RAN 5G SA network
Orange & Vodafone, Romania	2023	Samsung (vRAN), Wind River (abstraction layer), Dell (servers)	Shared rural pilot near Bucharest; 4G live Oct 2023; 2G/5G to follow	*Partial*	First live 4G calls over a shared Open RAN pilot in Europe; Samsung-led RAN stack; first virtualised 2G in Europe (May 2024 expansion)
Orange Spain	2023	Mavenir (RAN SW), Casa Systems (5GC), HPE, Dell	Lab/field pilot integrating Open RAN nodes with cloud-native 5G SA core (Oct 2023)	*Partial*	Sub-one-hour network activation via automation; multi-vendor RAN/core combination; built on Orange's Pikeo experimental 5G SA platform
Vodafone UK	2023	Samsung (vRAN & radios), Dell (servers), Wind River, Intel, Keysight	Replacement of 2,500 Huawei 4G/5G sites in Wales & SW England; multi-year	Island of interoperability	Samsung-led vRAN; O-RAN fronthaul compliant; KPIs exceed legacy Huawei in 4G/5G call success, up/download speeds
Vodafone Germany	2023	Samsung (vRAN & radios)	Commercial pilot at rural sites in SE Bavaria & NE Lower Saxony; start early 2023	Island of interoperability	First O-RAN MoU-compliant pilot in Germany; Samsung sole RAN supplier; foundation for Vodafone's 30% European O-RAN target by 2030
A1 Group, Bulgaria	2023	Nokia (Cloud RAN vDU, AirFrame servers)	5G SA Cloud RAN trial in Sofia; end-to-end L3 data call	Island of interoperability	Nokia sole supplier (vDU + servers + core); vDU served radio cells over 15 km fronthaul; validated cloud deployment automation
1&1 Drillisch, Germany	2022	Rakuten Symphony (end-to-end integrator), Mavenir (core)	Nationwide 5G greenfield; FWA launched Dec 2022; mobile (eMBB) launched Dec 2023	Island of interoperability	Europe's first fully virtualised Open RAN mobile network; single-integrator model; Rakuten Symphony responsible for end-to-end performance
Vodafone Idea, India	2023	Mavenir (end-to-end: CU, DU, OpenBeam radios, RAN SW on Red Hat OpenShift)	Commercial-phase pilot; live traffic on N78 (3.5 GHz) & N258 (mmWave); commenced Sep 2023	Island of interoperability	First O-RAN-compliant deployment in Vi's network; Mavenir sole end-to-end supplier; ahead of planned large-scale 5G rollout
Deutsche Telekom, Germany	2025	Nokia (vDU, O-RUs), 1Finity/Fujitsu (O-RUs), Mavenir (DU/CU SW, European footprint)	Brownfield replacement of >3,000 Huawei 5G sites in Northern Germany; commercial service live in Neubrandenburg area; completion targeted 2027	*Partial*	First large-scale Huawei replacement via Open RAN in Europe; Nokia paired with 1Finity radios over O-RAN open fronthaul; KPIs "on par with and better than" the incumbent; 30,000-site pan-European RFQ issued Jan 2026
TELUS, Canada	2025	Samsung (vRAN 3.0, Massive MIMO radios, RIC/SMO), Amplitech (FDD O-RUs), Wind River (cloud infra), HPE (DU servers, Intel vRAN Boost)	Nationwide brownfield rollout; 18% of sites O-RAN compliant by end-2025; target 40–50% converted to vRAN by end-2027	*Partial*	Rare brownfield Tier-1 example of genuinely multi-vendor Open RAN; Samsung baseband paired with third-party radios; operator-led systems integration (no neutral integrator); Samsung AI-powered RIC deployed Sep 2025
Viettel, Vietnam	2025	Viettel High Tech (in-house gNodeB: DU + RU), Qualcomm (X100 vDU SoC, QRU100 Massive MIMO RU platform, Edgewise Suite)	Commercial launch Nov 2024; >300 sites across Vietnamese provinces by Q1 2025; expansion to >2,000 sites planned by end-2025	Island of interoperability	World's first O-RAN 5G network using Qualcomm 5G RAN chipset platforms; fully in-house RAN stack developed by operator; "Make in Vietnam" industrial policy driver; exported to Middle East markets

emergence of "islands of interoperability," calling instead for neutral standards and certification programs to guarantee broad compatibility.

C. Technical and Market Realities

Despite these ambitions, full multi-vendor interoperability remains difficult to achieve. Large public test campaigns such as the U.S. NTIA 5G Challenge revealed that end-to-end multi-vendor operation requires extensive coordination, tuned configurations, and iterative fixes rather than automatic interchangeability [7]. Operator reports likewise describe non-trivial O-RU/O-DU bring-up and the need for sustained information sharing across vendors to complete integrations. Certification programs reflect this reality: the O-RAN ecosystem issues conformance certificates and *interoperability (IOT) badges* only for *specific pairs or systems*, not for all possible combinations, signaling that vendors and operators prioritize tested partner sets over universal compatibility [8]. Deutsche Telekom notes that true plug-and-play on the R1 interface for rApps is "not yet achieved," calling for stronger pre-integration and lab certification to reduce multi-vendor effort [9]. Global PlugFests and OTIC/i14y lab programs are accelerating maturity but still report "tackling issues" and expanding the set of validated combinations rather than declaring ubiquitous interchangeability [10].

The market evidence mirrors these technical constraints. Industry observers have coined the term "single-vendor Open RAN" to describe deployments where interfaces follow O-RAN specifications but all major components still come from one vendor. AT&T's nationwide rollout using $14 billion of Ericsson equipment is a prominent example: although the interfaces are open, the network remains vendor-homogeneous while six years into the Open RAN movement, there are barely any multi-vendor networks meeting full performance needs. Early commercial deployments therefore focus on less demanding scenarios, rural coverage, indoor networks, or greenfield operators, where the risk of lower initial performance is manageable.

Even when multiple suppliers are involved, operators and vendors gravitate toward known-good combinations validated in labs, forming *islands of interoperability* in which equipment works reliably within a cluster but not necessarily across clusters. European carriers have identified reducing integration complexity through better certification and testing as a prerequisite for expanding Open RAN to dense urban macro networks [11].

D. Graph Construction, Data Curation, and Methodological Limitations

To support the qualitative discussion with empirical evidence, we constructed a weighted co-occurrence graph of Open RAN ecosystem relationships. The graph was derived from publicly reported deployments, pilots, trials, plugfests, and ecosystem announcements. Each node represents an Open RAN actor, such as an operator, vendor, system integrator, or cloud provider. An undirected edge connects two actors when they are reported as participating in the same Open RAN activity, and the edge weight captures the frequency of such co-occurrences.

Before constructing the graph, actor names were normalized to reduce duplicate representations of the same entity, incomplete or ambiguous records were removed, and duplicate company pairs were aggregated by summing their weights. Company pairs were treated as unordered because the analysis focuses on co-occurrence rather than directionality. Isolated single-node components were removed so that the analysis focused only on observed relationships among actors.

The resulting graph is defined as $G = (V, E, w)$, where V is the set of Open RAN actors, E is the set of observed co-occurrence relationships, and w_{ij} is the edge weight between actors i and j. For each node, we computed its weighted degree, or node strength, as $s_i = \sum_{j \in N(i)} w_{ij}$, which indicates the relative visibility of an actor within the observed collaboration network.

To identify collaboration clusters, we applied weighted Louvain community detection. The detected communities are interpreted as empirical indicators of recurring Open RAN collaboration patterns and potential "interoperability islands." They do not prove formal lock-in or complete technical incompatibility, but they reveal repeated patterns of pre-validated or frequently reported collaboration.

The dataset includes only records that identify at least two concrete Open RAN actors and refer to a specific deployment, trial, plugfest, integration activity, or ecosystem collaboration. Generic marketing claims, duplicate entries, incomplete records, and relationships that could not be clearly inferred were excluded or consolidated to avoid double counting.

This methodology has limitations. The graph reflects publicly visible activity and may omit private trials, confidential evaluations, or unpublished integrations. Public announcements may also overrepresent successful collaborations. Moreover, co-occurrence does not necessarily imply deep technical interoperability, since actors may participate in the same deployment through a prime vendor, system integrator, or limited interface scope. Edge weights therefore capture observed collaboration frequency rather than deployment scale, traffic volume, performance, or commercial maturity. The graph should be interpreted as a snapshot of an evolving ecosystem, complementing the technical and qualitative analysis presented in the rest of the paper.

Scope of longitudinal and correlation analysis: A full longitudinal analysis of cluster evolution is outside the scope of this paper. Open RAN pilots, trials, plugfests, and early commercial deployments are still at a relatively early stage, while public records are sparse, irregularly reported, and not consistently timestamped. Therefore, the current dataset is best interpreted as a snapshot of observed ecosystem relationships rather than as a basis for robust temporal modeling. Similarly, correlations between deployment type and technical characteristics would require more standardized information on deployment date, interface scope, vendor role, certification status, and commercial maturity. Such analysis would be valuable in future work as Open RAN deployments mature and richer datasets become available.

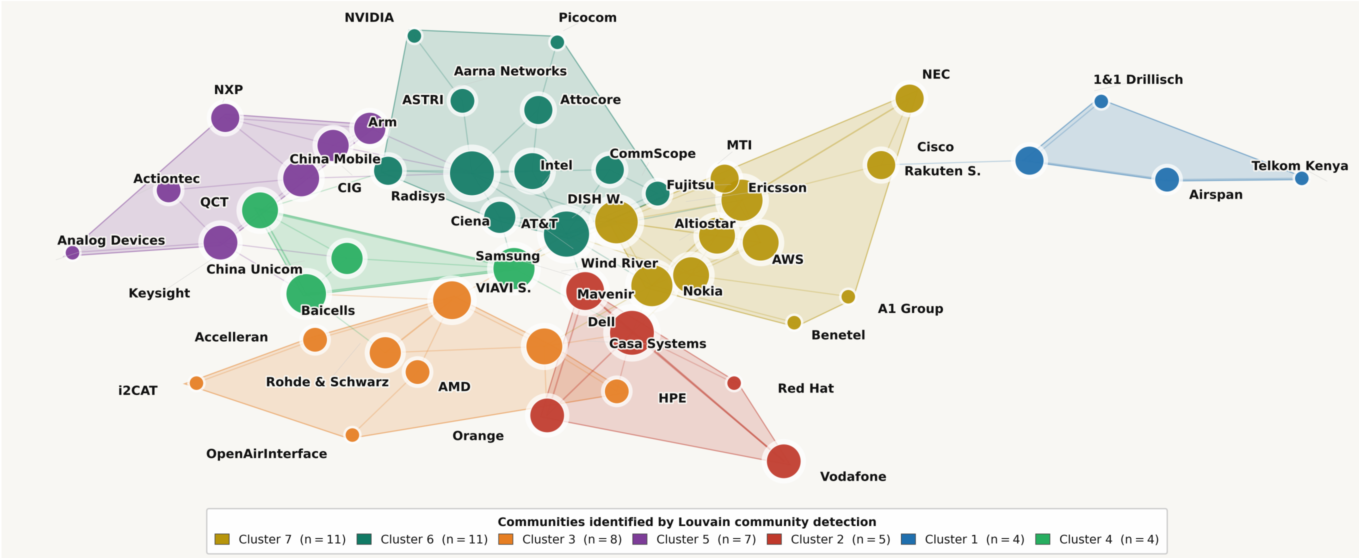

Fig. 2. Company co-occurrence graph derived from 31 Open RAN activities (9 field deployments and 22 plugfest/lab demonstrations, 2020–2025). Nodes represent the 50 participating vendors and operators; edges connect companies that co-appeared in the same activity, with edge weight proportional to repeated co-participation (97 edges total). Node size encodes weighted co-participation degree (strength). Seven communities are identified by Louvain community detection (resolution $= 1.5$; modularity $Q = 0.575$); community sizes are 11, 11, 8, 7, 5, 4, and 4 companies respectively. Convex hulls delineate each community. The seven clusters correspond to the "islands of interoperability" discussed in the text.

E. Technical Challenges of Disaggregated Architectures

While the Open RAN vision is compelling, disaggregating the RAN stack introduces a distinct class of technical challenges that do not arise, or arise in much milder form, in vertically integrated systems. We group these challenges into five categories.

Interface Multiplicity and Integration Overhead. A single Open RAN site requires the correct operation of at least six distinct inter-domain interfaces: the Open Fronthaul (O-RU to O-DU), F1 midhaul (O-DU to O-CU), E2 (near-RT RIC to RAN nodes), A1 (non-RT to near-RT RIC), O1 (SMO to network functions), and O2 (SMO to cloud infrastructure). Each interface is specified by the O-RAN Alliance, but specification compliance does not guarantee seamless interoperability: vendors interpret optional fields differently, implement subsets of mandatory procedures, and expose vendor-specific configuration parameters with no standardised defaults. First-hand observations from private 5G testbeds, including SONIC Labs and Vodafone's validation facilities, confirm that even with well-defined F1, E2, and A1 interfaces, integration is routinely hindered by inconsistent vendor implementations, significant manual configuration overhead, and the near-total absence of automated CI/CD-based testing pipelines that could catch regressions across vendor software updates [5]. The O-RAN Alliance acknowledges this reality through its network of Open Testing and Integration Centres (OTICs), which provide common laboratories for functional verification and integration badging, but the existence of more than fifteen such centres worldwide signals that integration effort is substantial rather than incidental [6].

Real-Time Constraints Across Vendor Boundaries. The RAN is fundamentally a hard real-time system. Scheduling decisions, beamforming weight updates, and handover executions must occur within microsecond-to-millisecond windows that leave no margin for additional latency introduced by inter-vendor communication overhead. In a single-vendor stack, these interactions occur over proprietary internal buses with deterministic timing; in a disaggregated stack, they traverse standardised interfaces whose latency budgets are tight but whose implementations vary. The Open Fronthaul, for instance, imposes strict timing constraints on the transport network connecting the O-RU to the O-DU, and meeting these constraints with third-party hardware combinations requires careful link engineering that eliminates the simplicity that operators expect from plug-and-play deployment. The scale of this challenge was made concrete by the U.S. NTIA 5G Challenge, which demonstrated that end-to-end multi-vendor operation requires extensive coordination, iterative tuning, and repeated integration fixes rather than the automatic interchangeability implied by the standard [7].

Performance Parity Gap. As a direct consequence of integration complexity and real-time overhead, multi-vendor Open RAN deployments have not yet achieved consistent performance parity with single-vendor systems in the most demanding network scenarios. This explains a clearly observable pattern in early deployments: commercial Open RAN rollouts are disproportionately concentrated in rural coverage extensions, indoor enterprise networks, private 5G deployments, and greenfield operators, precisely the scenarios where traffic density is lower, performance targets are less stringent, and the risk of falling short of single-vendor baselines is manageable [11]. Full-scale nationwide deployment in dense urban macro networks, where spectral efficiency, peak throughput, and handover success rates are under constant pressure, remains a future milestone rather than a current achievement. The industry shorthand “single-vendor Open RAN”, where interfaces follow O-RAN specifications but all major components come from one supplier, has emerged as a transitional form that preserves performance while deferring the harder problem of cross-vendor integration, as illustrated by AT&T’s $14 billion Ericsson deployment.

xApp Conflict and RIC Stability. The RAN Intelligent Controller introduces a new category of technical risk that has no equivalent in traditional RAN: the possibility of destabilising feedback loops caused by concurrent control applications. When multiple xApps act simultaneously on overlapping control levers, such as handover offset adjustments, physical resource block (PRB) quota assignments, or radio sleep-mode scheduling, their individual actions can interact in ways that none of them was designed to handle, producing oscillatory behaviour, resource starvation, or performance degradation that is difficult to attribute to any single application [12]. This challenge is compounded by the open nature of the RIC platform: third-party xApps are developed independently, may not disclose their internal logic, and cannot be assumed to coordinate with one another. Mitigations discussed in the literature and at the WCNC 2025 Open RAN Workshop include scoped API boundaries that limit each xApp’s control domain, lightweight priority and preemption guards, canary and A/B rollout mechanisms to detect unstable behaviour before fleet-wide deployment, and a minimal coordination layer to resolve conflicts before they propagate to the RAN data path [1]. These are active areas of standardisation and research, not solved problems.

Certification Gap. The O-RAN Alliance’s conformance and interoperability testing regime, while well-intentioned and growing, is structurally misaligned with the goal of universal plug-and-play compatibility. Conformance certificates and Interoperability (IOT) badges are issued for *specific vendor pairs or pre-validated system combinations*, not for all mathematically possible multi-vendor permutations [8]. This means that an O-RU certified to interoperate with O-DU A may not interoperate correctly with O-DU B, even if both O-DUs hold valid conformance certificates. Deutsche Telekom has publicly stated that true plug-and-play on the R1 interface for rApps is “not yet achieved,” calling for stronger pre-integration and lab certification to reduce the multi-vendor integration burden [9]. Global PlugFests and OTIC programmes are progressively expanding the set of validated combinations, but the current state is one of *expanding test matrices* rather than *declared universal interchangeability* [10]. Until certification transitions from pair-wise to combinatorial coverage, or until automated conformance testing can scale to cover a much larger portion of the combination space, the certification gap will remain a structural barrier to the “any-vendor-anywhere” interoperability that the Open RAN vision promises.

Integration challenges. The emergence of interoperability clusters can also be explained by concrete Open RAN integration challenges. For example, O-RU/O-DU integration depends not only on the open fronthaul interface, but also on timing, synchronization, radio feature support, conformance profiles, and vendor-specific implementation choices [11], [13], [14]. Similarly, Near-RT RIC integration depends on E2 service-model compatibility, xApp control logic, and conflict management when multiple xApps act on overlapping RAN parameters [1], [12]. In the management domain, SMO/O1 integration may require adaptation across different configuration models, telemetry formats, and fault-management behaviours [1], [6]. These examples explain why interoperability is often validated for specific vendor combinations through testing, certification, and plugfests, rather than guaranteed across all theoretically compatible components [7], [9], [15], [16]. Where possible, such impacts could be quantified through integration effort, certification scope, number of validated vendor combinations, conformance results, or operational indicators such as latency, handover performance, energy consumption, and fault-resolution time. However, these measurements are not consistently reported in public Open RAN announcements.

IV. Global Case Studies of Open RAN Deployment

Real-world pilots and commercial rollouts provide insight into the current maturity of 5G [17] and Open RAN technologies. To quantify the worldwide status of Open RAN development, we curated and analysed **31 activities** , comprising 9 commercial or pre-commercial field deployments (3 greenfield, 6 brownfield) and 22 plugfests, lab integrations, and conference demonstrations , drawn from public operator announcements, O-RAN Alliance PlugFest records, and MWC exhibition catalogues spanning 2020–2025. The 31 activities span four regions: **Europe** (20), **North America** (8), **Asia-Pacific** (3), and the Middle East / Africa / South-East Asia corridor (not yet represented in this curated set, though emerging deployments in Kenya, Kuwait, Vietnam, and Bangladesh are noted in extended operator data). Together they involve **50 unique vendors and operators**.

From each activity we extracted the participating companies and built a company co-occurrence graph: nodes are firms; edges connect firms that co-appeared in the same activity, with weights increasing for repeated collaborations. Node size encodes each firm’s co-participation intensity (weighted degree). Communities are identified with Louvain community detection (resolution $= 1.0$, $Q = 0.575$).

Applying the taxonomy introduced in Section II to the 9 field deployments (Table I), we find that **2 deployments qualify as fully multi-vendor** (Rakuten Mobile Japan and DISH Wireless USA, both greenfield operators with 4–6 distinct RAN vendors), **2 are partial** (Orange Romania and Orange Spain, where multi-vendor scope is confined to the cloud/IT layer while the RAN itself remains single-supplier), and **5 are islands of interoperability** (Vodafone UK, Vodafone Germany, A1 Bulgaria, 1&1 Germany, and Vodafone Idea India), each dominated by a single RAN supplier operating within O-RAN-compliant interfaces. This 2-2-5 split , only 22% of field deployments achieving true multi-vendor RAN diversity , provides a concrete quantitative measure of how far the ecosystem remains from the original Open RAN vision.

As shown in Fig. 2, seven distinct “islands of interoperability” already emerge from the co-occurrence data. This is a curated set of major, well-documented activities and is not exhaustive, but the cluster structure confirms that collaboration patterns in the real ecosystem are far from random: companies repeatedly co-deploy with the same small set of partners, solidifying the island topology.

V. Market Implications: Vendors, Operators, and the Evolving Supply Chain

Although Open RAN remains in an early and partly fragmented stage, its influence on the telecommunications ecosystem is already significant. The technology is reshaping competitive dynamics among incumbent suppliers, enabling the emergence of new entrants, altering operator procurement strategies, and drawing information technology (IT) vendors more deeply into the radio access network (RAN) supply chain.

A. Incumbent Vendors

For the three dominant RAN suppliers, Ericsson, Nokia, and Huawei, Open RAN presents both a competitive threat and a catalyst for strategic adaptation. Huawei, excluded from many Western markets, remains outside the O-RAN Alliance and faces declining market share in Europe as political pressures constrain its reach. Ericsson participates in O-RAN standardization but continues to promote the reliability of its integrated solutions. Nokia, by contrast, has pursued an aggressive “anyRAN” strategy, declaring O-RAN compliance across its 5G AirScale portfolio and demonstrating interoperability with third-party radios. This positioning allows operators to adopt openness while retaining Tier-1 vendors for key functions such as the O-DU and O-CU. All major incumbents are simultaneously developing Cloud RAN offerings to remain competitive as networks become increasingly software-defined.

B. Emergent Vendors

A growing cohort of specialized suppliers has entered the market to exploit opportunities in the disaggregated RAN value chain. Software-centric firms such as Mavenir, Parallel Wireless, and Altiostar/Rakuten Symphony focus on virtualized baseband and RIC software, while hardware specialists such as Benetel and ASOCS concentrate on O-RAN–compliant radio units and accelerators. Open RAN lowers barriers to entry by allowing newcomers to compete in specific functional domains rather than delivering an entire RAN stack. For example, Benetel has supplied radios for private 5G trials in Europe, including a Polish pilot in which a local software vendor (IS-Wireless) provided the RAN software integrated with T-Mobile’s core network. Nevertheless, these challengers face scale and financial constraints: developing carrier-grade RAN software remains capital-intensive, and many projects remain in the trial phase without large revenue streams. Mavenir, for instance, postponed its initial public offering in 2020 and reduced staff following slower-than-expected uptake in key markets.

C. Mobile Network Operators

Mobile network operators stand to benefit from increased supplier diversity but also assume new technical and organizational risks. In theory, broader vendor choice enhances competition and lowers costs, and operators frequently leverage Open RAN commitments to negotiate more favorable terms with incumbents. In practice, however, operators adopting Open RAN often become de facto system integrators, requiring in-house expertise or third-party services to integrate multi-vendor networks. Integration complexity can erode the anticipated cost savings and exacerbate challenges related to energy efficiency, feature parity, and performance optimization. Nonetheless, leading operators, including Vodafone, Orange, Telefónica, Rakuten, Dish, Verizon, AT&T, and BT, have publicly committed to open architectures. Geopolitical factors reinforce this trend: for example, UK operators replacing Huawei equipment are using Open RAN to introduce new suppliers such as Samsung, with Vodafone’s UK rollout replacing 2,500 Huawei sites with Samsung-based Open RAN gear. Such initiatives also align with government-backed programs in the UK, EU, Japan, and the United States that fund Open RAN testbeds and supplier diversification.

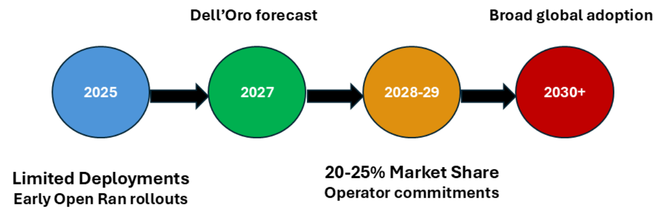

Fig. 3. Projected Open RAN market adoption timeline from early deployments to broader global uptake.

D. Broader Supply Chain and IT Convergence

Open RAN blurs the boundaries between telecommunications and information technology. Virtualized RAN functions create opportunities for semiconductor firms (Intel, Marvell, Qualcomm) and server vendors (Dell, HPE), as demonstrated by deployments where Nokia baseband software runs on Dell PowerEdge servers with Intel processors or Orange pilots operate on HPE cloud infrastructure. Cloud and software companies such as Microsoft, VMware, and Amazon are also positioning themselves to provide telco cloud platforms and RAN automation frameworks. This diversification increases ecosystem resilience by reducing dependence on a small set of traditional vendors, but it also expands the attack surface and operational complexity, reinforcing the need for rigorous certification, interoperability testing, and zero-trust security practices [18].

VI. Future Outlook and Recommendations for Open RAN

Recent bibliometric and thematic analyses confirm that Open RAN research is transitioning from conceptual exploration to mature, deployment-focused studies. Singh and Samal [19] conducted a comprehensive review of global Open RAN publications, identifying key trends such as increased collaboration between telecom vendors and cloud providers, the rise of AI-driven RIC innovations, and a growing emphasis on sustainability and network intelligence. Open RAN is often described as a journey rather than a destination. As we look to the future, there is broad consensus that Open RAN will play an increasingly significant role in mobile networks, but also that full realization of its benefits will take time. As summarized in Fig. 4, the Open RAN trajectory is staged across 2025–2035, outlining short-, mid-, and long-term milestones, enabling technologies, and persistent challenges that collectively frame a pragmatic adoption pathway.

Gradual Expansion and Market Share: Forecasts by industry analysts indicate a steady rise in Open RAN adoption over the next decade. Dell'Oro Group, for example, predicts that by 2028–2029, over 25% of the global RAN market (by revenue) could be served by Open RAN architectures – but importantly, this figure includes single-vendor O-RAN deployments.

Truly multi-vendor RAN (mix-and-match) is expected to be a smaller subset (Dell'Oro projects $\leq$ 10% of the total RAN market by 2029 might be multi-vendor interoperable in the fullest sense). In other words, Open RAN is likely to succeed in general, to the point that many networks will use O-RAN compatible equipment (perhaps even by default), but the purist vision of many vendors all competing at every site will materialize more slowly. This nuance is important for setting expectations. These market projections should be interpreted as indicative estimates rather than deterministic predictions. Open RAN adoption remains dependent on operator procurement strategies, certification maturity, integration costs, vendor participation, macroeconomic conditions, and the pace of standardization. In particular, the distinction between single-vendor O-RAN-compliant deployments and genuinely multi-vendor interoperable deployments introduces additional uncertainty when interpreting market-share forecasts.

Near-Term Deployments – Targeted Use Cases: In the next 2–3 years, we expect to see Open RAN deployments focused on specific scenarios e.g., rural coverage, indoor enterprise networks, private 5G networks, and network extension in less dense areas. These are domains where Open RAN solutions are already proving themselves (as seen in case studies) and where the risk of using a new approach is lower. By 2025, Europe's major operators anticipate moving from trials to initial commercial deployments in urban or suburban areas as well, but likely still with some constraints (e.g. maybe one macro layer or one city quarter with Open RAN, before wide-scale citywide rollout). Full-scale nationwide Open RAN in dense urban networks will probably come later in the decade, once performance and energy efficiency are clearly on par. It's worth noting that improving energy efficiency and 5G feature parity for Open RAN is a top R&D priority now – we can expect advancements in those areas soon (for instance, better hardware accelerators for vRAN to reduce power draw, smarter sleep modes for radios, etc.).

Integration of AI and RIC: A promising aspect of Open RAN is the RAN Intelligent Controller (RIC), which provides a platform for running third-party apps (xApps/rApps) to optimize the network (for example, dynamic spectrum sharing, traffic steering, energy savings). In the future, Open RAN architectures can unleash more innovation through these software applications. We foresee a growing ecosystem of specialized RIC application developers – companies that are not radio vendors per se but offer AI-driven algorithms to improve network performance. Opening interfaces at the higher layers (like the RIC APIs) enables a sort of "app store" for network features. While still nascent, trials of RIC (e.g., by Dish, Vodafone, Telefónica) have shown the potential to, say, improve handover decisions or detect anomalies in real time. In a few years, operators could routinely deploy third-party xApps to add capabilities to their network without waiting for a monolithic vendor upgrade.

Stronger Certification and Testing Frameworks: To address the interoperability island issue, the next few years will likely see the maturing of certification labs and processes. We may have a situation akin to Wi-Fi certification in the past – a rigorous suite of tests that vendors must pass so that operators trust multi-vendor gear. The O-RAN Alliance and TIP are working on this (with initiatives like the OTIC labs – Open Testing and Integration Centres). Governments are also funding test facilities (the UK's SONIC lab, NTIA's 5G Challenge in the US, etc.) to facilitate multi-vendor interop testing. By establishing a common bar of interoperability (perhaps even an "Open RAN Certified" stamp), operators can mix vendors with less fear. Plugfests will continue and expand, possibly focusing on specific scenarios (massive MIMO interop, 5G SA features, etc.). The end goal is to make multi-vendor deployments almost as plug-and-play as single-vendor, which will significantly lower the integration cost barrier.

Security and Reliability Focus: As Open RAN becomes more prevalent, there will be heightened focus on security. Disaggregated networks mean a larger attack surface (more interfaces, more vendors). The future will see efforts to harden Open RAN security – for example, ensuring all O-RAN components comply with 3GPP and O-RAN security specifications (the EU's cybersecurity agency ENISA has already analyzed Open RAN security). The zero-trust approach will be emphasized, where each component is treated as potentially untrusted until verified. In practice, this could mean more robust authentication and encryption on O-RAN fronthaul links, vetting of software from new vendors, and perhaps guidelines from governments on approved Open RAN suppliers (similar to how core network vendors are vetted). Reliability will also be key – operators will demand that an Open RAN network can achieve the "five 9s" reliability of traditional networks. This might lead to more redundancy solutions (e.g., geo-redundant DUs, etc.) and service management standards for multi-vendor environments.

Economic and Policy Drivers: The economics of Open RAN are expected to improve over time. As the ecosystem matures, competition could drive down unit costs of hardware (e.g., remote radio units) and software licensing. Moreover, using COTS (commercial off-the-shelf) hardware and cloud infrastructure could yield OPEX savings at scale (especially if automation reduces manual operations). However, one should not expect dramatic cost savings immediately as the true measure of Open RAN's success will be cost reductions over a long term through increased competition and virtualization efficiencies. Policymakers, for their part, are likely to continue encouraging Open RAN deployment. We might see more targeted funding or incentives for operators to deploy Open

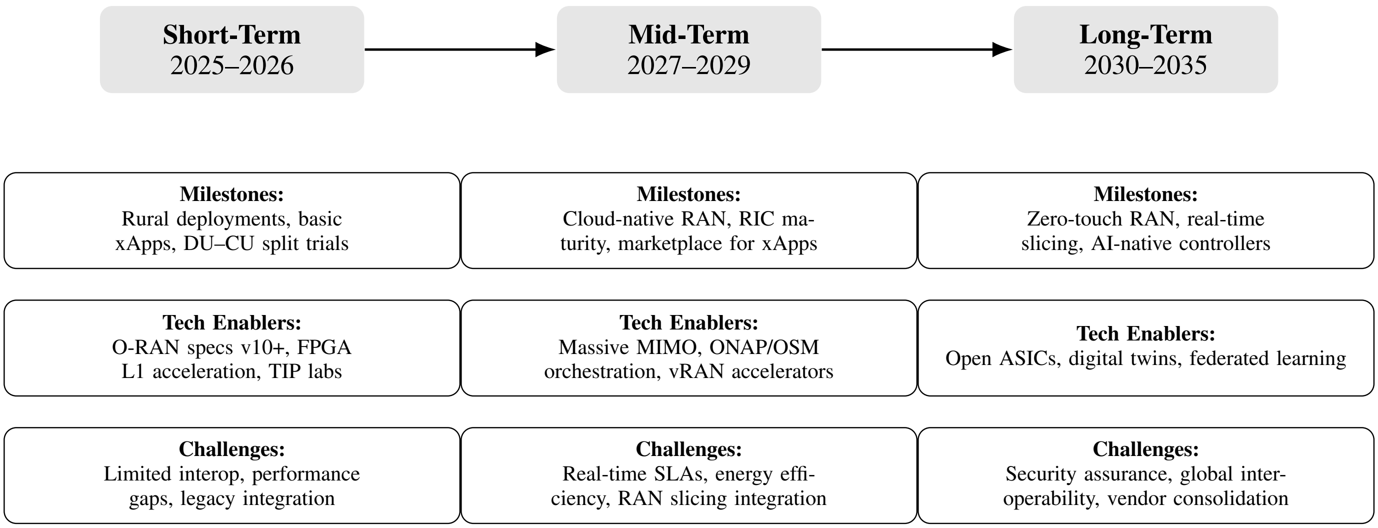

Fig. 4. Open RAN evolution roadmap across short-, mid-, and long-term phases. The figure summarizes key milestones, enabling technologies, and technical challenges from early rural deployments and basic xApps to cloud-native RAN, RIC maturity, zero-touch automation, AI-native controllers, and global interoperability.

RAN in underserved areas or to work with local suppliers. Japan, for instance, subsidized operators for open network equipment procurement; the US has allocated funding for Open RAN trials. This supportive policy environment could accelerate adoption, albeit with the caveat that technology readiness dictates the pace to some extent.

Global Spread and New Entrants: By 2025–2026, expect more countries to have at least one commercial Open RAN deployment. We may see brownfield operators in Europe and Asia slowly extend Open RAN from rural into urban. New 5G entrants (in countries introducing a fourth operator, for example) will almost certainly choose Open RAN to reduce upfront costs – following the Rakuten/Dish playbook. Additionally, developing markets might leverage Open RAN for cost-effective rural coverage (e.g., pilots in Africa or Southeast Asia with Open RAN for 4G/5G in rural communities). Organizations like the O-RAN Alliance now have over 300 members, including many from emerging markets, so the interest is worldwide. As the ecosystem grows, some consolidation may happen: not all new vendors will survive, but those with unique value (be it an innovative radio design, a great scheduling algorithm, etc.) could be acquired by larger firms or might carve out a sustainable niche.

A. Recommendations

For Open RAN to reach its potential, a few recommendations can be made for stakeholders:

For Operators: Remain engaged in trials and contribute learnings to the community. Invest in developing in-house integration skills or partner closely with integrators. Start with small-scale deployments to gain confidence, then scale gradually. Insist on standards compliance in RFPs to keep vendors honest. Collaborate with peers (as seen in the Orange/Vodafone partnership) to share knowledge and even infrastructure where appropriate.

For Vendors: Embrace openness genuinely – implement the O-RAN specs fully and participate in plugfests. Optimize solutions (hardware or software) for power and performance to meet parity with incumbents. Differentiate through innovation (for example, an AI-powered RAN optimizer or an energy-efficient radio design) to provide value beyond just being "open." Be prepared to partner – even rivals may need to interoperate in an open environment.

For Policy Makers: Continue supporting a diverse ecosystem through funding testbeds, R&D (especially in areas like open RAN security, and power-efficient chip design for vRAN), and avoiding heavy-handed mandates that could force immature tech too fast. Instead, use targets as guiding lights (like the UK's 35% by 2030 goal) and back the industry with resources to get there. International cooperation (as seen in joint statements) is beneficial – align on Open RAN principles and security frameworks so that vendors face a unified set of requirements.

For Researchers and Academics: There are ample areas for research: algorithms for multi-vendor orchestration, AI in the RIC, advanced antenna techniques that are vendor-agnostic, security mechanisms for open interfaces, etc. Engaging with the O-RAN Alliance working groups or operator testbeds can ground research in real-world needs.

B. Validation, Generalizability, and Boundary Conditions

The emergence of "interoperability islands" should be interpreted as a pattern supported by the observed co-occurrence network and the qualitative evidence discussed in this paper, rather than as a universal property of all Open RAN deployments. The detected clusters indicate that Open RAN collaborations often form around recurring combinations of operators, vendors, integrators, and cloud providers. This supports the argument that practical interoperability may remain partly bounded by pre-tested vendor groupings, integration experience, certification scope, and deployment-specific engineering constraints.

The generalizability of this conclusion is strongest for publicly reported commercial deployments, trials, plugfests, and ecosystem partnerships where multi-vendor integration is visible. However, the conclusion may be less applicable in scenarios where an operator enforces strict vendor-neutral integration requirements, where open interfaces are fully validated across a broad set of vendors, or where a mature certification framework guarantees interoperability beyond pairwise combinations. It may also be weaker in private trials or internal operator laboratories that are not publicly documented.

This interpretation is consistent with prior discussions of Open RAN interoperability, which emphasize that open interfaces alone do not automatically ensure plug-and-play substitutability. Compared with studies that focus mainly on architectural openness or standardization progress, our analysis highlights the ecosystem-level structure of actual collaborations. The results therefore complement existing technical and standardization-oriented work by showing how practical Open RAN deployment patterns may still concentrate around recurring interoperability clusters.

VII. WCNC 2025 Panel Reflections

The industry panel in the Open-RAN Workshop at WCNC 2025 (Milan) offered a pragmatic assessment of how aspirations for "openness" intersect with operational constraints. Discussion moved beyond normative positions to focus on the loci of integration frictions observed in early deployments and the operational patterns emerging to manage them.

A. Performance Accountability in Multi-Vendor Settings

Panelists noted that the transition from vertically integrated stacks to disaggregated supply chains replaces a single, well-defined escalation path with a distributed responsibility model. Fault localization must traverse organizational and tooling boundaries; performance degradations frequently manifest at interface junctures (e.g., fronthaul timing budgets, E2 message semantics) rather than within a single component. Two recurring mitigations were highlighted: (i) *observability as a first-class concern*, standardized telemetry, correlatable traces, and self-contained, reproducible test bundles attached to support tickets; and (ii) *bounded pre-integration*, publication of validated "reference constellations" (explicit vendor pairings with version pins and service-level objectives). Consequently, many operators adopt a staged approach: begin with a curated coalition (occasionally a single-vendor Open RAN baseline), and expand the partner matrix as monitoring, tooling, and operational playbooks mature [7].

B. Interoperability of xApps and the Role of Coordination

While the RIC promises diversity of control logic, concurrent xApps acting on overlapping levers (e.g., handover offsets, PRB quotas, energy states) can induce contention and unstable feedback loops. The panel converged on treating the RIC as a *governed platform*: clearly scoped APIs and ownership of control domains; lightweight policy guards (priority, preemption, and admission control); safe-change mechanisms (sandboxes, canary/A–B rollouts, and circuit breakers); and a minimal coordination layer to resolve conflicts [12] before they affect the RAN data path. This framing preserves innovation at the network's control edge while maintaining an explicit stability envelope [1].

C. "Islands of Interoperability" as a Transitional State

Universal mix-and-match remains the long-term objective, but the near-term equilibrium resembles *certified clusters*. Within each cluster, interfaces are open, software versions are pinned, and combinations are exercised in labs and PlugFests; portability across clusters is improving but not yet automatic. Panelists characterized this not as a retreat from openness but as a workable staging model: expand clusters by widening the test matrix; bridge clusters through neutral conformance and IOT badging; and avoid re-creating closed gardens by publishing evidence artifacts (configs, traces, results) alongside every certified stack. Under this view, openness accrues incrementally, first within clusters, then across them, anchored in testable guarantees rather than declarative claims.

VIII. Conclusion

This paper examined the gap between the policy ambition of fully open, mix-and-match RANs and what early deployments currently deliver. From global case studies and a co-occurrence analysis of pilots and trials, we showed that the ecosystem is consolidating into *islands of interoperability*, curated, pre-validated constellations that perform reliably within bounded scopes. These islands are a pragmatic response to real-time constraints, integration risk, and commercial incentives, yet they also pose a material threat to the promise of openness championed by governments and industry coalitions: if allowed to harden, they can become de-facto closed ecosystems that simply relocate vendor lock-in. At the same time, Open RAN's core proposition, standards-based disaggregation, increased supplier diversity, and a programmable control plane via the RIC enabling AI/ML-driven optimization, has delivered tangible progress; openness is advancing, but in staged, evidence-backed steps. The synthesis of these findings is twofold: (i) today's openness is real but qualified, realized primarily within certified multi-vendor stacks; and (ii) avoiding a relapse into re-lock-in requires explicit countermeasures, performance-realistic certification that rewards cross-cluster portability, common operational and security practices across interfaces, and governed RIC programmability that enables third-party intelligence without destabilizing the data path. With such guardrails, curated ecosystems can serve as transitional scaffolding rather than endpoints, converting islands into bridges and aligning practical deployment with the original policy intent of broad, durable, and economically meaningful openness. The road to a fully open RAN is therefore evolutionary rather than binary: progress will be incremental, balancing interoperability with rock-solid performance, but, if the community sustains collaborative engineering and rigorous validation, open enough to materially transform mobile networks.

REFERENCES

[1] K. Alam, M. A. Habibi, M. Tammen, D. Krummacker, W. Saad, M. D. Renzo, T. Melodia, X. Costa-Pérez, M. Debbah, A. Dutta, and H. D. Schotten, "A comprehensive tutorial and survey of o-ran: Exploring slicing-aware architecture, deployment options, use cases, and challenges," *IEEE Communications Surveys & Tutorials*, vol. 28, pp. 1637–1678, 2026. doi: 10.1109/COMST.2025.3598406 .

[2] Y. I. H. Aljanabi, S. Y. Hussain, P. T. Mamataevna, V. S. Al-Doori, and J. M. Brieg, "Advancements in open ran and the decentralization of telecom networks," *Journal of Information Processing and Management*, vol. Special Issue, pp. 1213–1245, 2025. doi: 10.22034/jipm.2025.728406 .

[3] "Ieee standard glossary of software engineering terminology," *IEEE Std 610.12-1990*, pp. 1–84, 1990. doi: 10.1109/IEEESTD.1990.101064 .

[4] A. S. Abdalla, P. S. Upadhyaya, V. K. Shah, and V. Marojevic, "Toward next generation open radio access networks: What o-ran can and cannot do!" *IEEE Network*, vol. 36, no. 6, pp. 206–213, 2022. doi: 10.1109/MNET.108.2100659 .

[5] X. Krasniqi, E. Hajrizi, and B. Qehaja, "Challenges and lessons learned during private 5g open ran deployments," in *Proc. Int. Conf. on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*, 2023, pp. 1–7. doi: 10.1109/ICECCME57830.2023.10252662 .

[6] V. Yanover, *Vendor Interoperability, System Integration, and Maintenance*. John Wiley & Sons, 2025, pp. 83–87. doi: 10.1002/9781394241170.ch4 .

[7] M. H. Pinson, M. Poletti, J. E. Kub, J. Glenn, S. T. Thompson, R. Kupsh, N. Areqat, O. Dharmadhikari, A. George, T. L. Hardin, C. Johnson, S. Kapoulas, and S. Sriram, "5g challenge preliminary event: Evaluating modular, interoperable, multi-vendor, open ran solutions," U.S. Department of Commerce, National Telecommunications and Information Administration, Institute for Telecommunication Sciences, Technical Memorandum NTIA TM-23-568, May 2023. doi: 10.70220/nlkk7f2g . [Online]. Available: https://doi.org/10.70220/nlkk7f2g

[8] J. Groen, S. D'Oro, U. Demir, L. Bonati, M. Polese, T. Melodia, and K. Chowdhury, "Implementing and evaluating security in o-ran: Interfaces, intelligence, and platforms," *IEEE Network*, vol. 39, no. 1, pp. 227–234, Jan 2025. doi: 10.1109/MNET.2024.3434419 .

[9] F. Mehran, C. Turyagyenda, and D. Kaleshi, "Experimental evaluation of multi-vendor 5g open rans: Promises, challenges, and lessons learned," *IEEE Access*, vol. 12, pp. 152 241–152 261, 2024. doi: 10.1109/ACCESS.2024.3476963 .

[10] I. Wong, *Open RAN Test and Integration*, 2024, pp. 172–190. doi: 10.1002/9781119886020.ch11 .

[11] T. V. Ngo, M. V. Ngo, B. Chen, G. Gemmi, E. Baena, M. Polese, T. Melodia, W. Chien, and T. Quek, "Consistent and repeatable testing of o-ran distributed unit (o-du) across continents," in *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, 2024, pp. 1–5. doi: 10.1109/VTC2024-Fall63153.2024.10757917 .

[12] A. Wadud, F. Golpayegani, and N. Afraz, "Qacm: Qos-aware xapp conflict mitigation in open ran," *IEEE Transactions on Green Communications and Networking*, vol. 8, no. 3, pp. 978–993, 2024. doi: 10.1109/TGCN.2024.3431945 .

[13] Vodafone and N. DOCOMO, "O-ru/o-du integration challenges in an open ran," *White Paper*, 2023, describes the work needed to integrate radio units and distributed units from different vendors.

[14] O-RAN ALLIANCE WG5, "O-ran interoperability test specification, version 11.0," Tech. Rep., 2024.

[15] O-RAN ALLIANCE, "O-ran certification and badging program," 2023, available at https://www.o-ran.org/certification-badging.

[16] ——, "Overview of open testing and integration centre (otic) and o-ran certification and badging program," 2023, white Paper, accessed Jan. 2025.

[17] C. A. Gutierrez, O. Caicedo, and D. U. Campos-Delgado, "5g and beyond: Past, present and future of the mobile communications," *IEEE Latin America Transactions*, vol. 19, no. 10, pp. 1702–1736, 2021. doi: 10.1109/TLA.2021.9477273 .

[18] A. Kord, J. B. Coder, and V. Le, "Evolving open ran interoperability: A large-scale definition," in *2024 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2024, pp. 2095–2100. doi: 10.1109/ICCWorkshops59551.2024.10615745 .

[19] S. Singh and U. Samal, "Insights and trends in open ran: The future of mobile networks," *Journal of Network and Systems Management*, vol. 33, no. 1, pp. 1–22, 2025. doi: 10.1007/s10922-025-09920-5 .

Nima Afraz is a Tenured Assistant Professor in the School of Computer Science at University College Dublin and leads the UCD Network Intelligence and Automation Lab (NIA). His research focuses on Open RAN, blockchain applications in telecommunications, network automation, network economics, and infrastructure virtualisation. He is a funded investigator at the Research Ireland CONNECT Centre and has secured nearly €2 million in research funding through major national and European grants, including the National Space Subsystems and Payloads Initiative, the Government of Ireland Postdoctoral Fellowship, and the EU-funded HORIZON MSCA Staff Exchange project RE-ROUTE. Dr. Afraz completed his PhD in Computer Science at Trinity College Dublin in 2020, where his work on multi-tenant optical access networks contributed to standardisation and policy research. His research has influenced industry and public-policy outcomes, including Broadband Forum TR-402, ETSI specifications, and World Bank work on affordable broadband deployment. He also serves as Vice-Chair of the Linux Foundation's Hyperledger Telecom Special Interest Group, supporting industry-academic collaboration in open-source telecommunications innovation.

Mohammad Shojafar (Senior Member, IEEE) received a PhD degree in Information and Communication Technology from Sapienza University of Rome, Italy, in 2016, where he was awarded the distinction of "Excellent" for his doctoral research. He is currently an Associate Professor of Network Security and an Intel Inno-vator with the 5G/6G Innovation Centre (5G/6GIC), Institute for Communication Systems, University of Surrey, U.K. Before this, he was a Senior Researcher and a Marie Curie Fellow with the SPRITZ Security and Privacy Research Group at the University of Padua, Italy. His research interests include 5G/6G network security and privacy, Open RAN security, and adversarial artificial intelligence for fog and cloud networking. Dr. Shojafar has secured approximately £2 million in research funding as Principal Investigator for various EU- and UK-funded projects. He serves as an Associate Editor for the IEEE Transactions on Network and Service Management, IEEE Transactions on Intelligent Transportation Systems, IEEE Transactions on Green Communications and Networking, and IEEE Transactions on Consumer Electronics. He is a Professional Member of the ACM, an ACM Distinguished Speaker, a Fellow of the Higher Education Academy (FHEA), and a Marie Curie Alumni.

Johann Marquez-Barja (Senior Member, IEEE) is Full Professor at the University of Antwerp and imec, Belgium, where he leads the Programmable & Intelligent Networks (PIN) group at IDLab. Ranked among the world's Top 2% scientists (Stanford/Elsevier), he has coordinated and contributed to numerous European research projects in 5G/6G, edge computing, network programmability, and vehicular communications. He also leads the Citylab Smart City, SmartHighway, and nextGOpen testbeds in Antwerp, Belgium. Prof. Marquez-Barja has authored over 300 publications, delivered several keynotes, and received 35+ awards. He is an IEEE Distinguished Lecturer (IEEE Vehicular Technology Society & IEEE Education Society), Senior Member of IEEE and ACM, and Fellow of EAI. His research spans 5G/6G architectures, network-of-networks, interoperability, and smart mobility.

Hamed Ahmadi (Senior Member, IEEE) received the Ph.D. degree from the National University of Singapore in 2012. He is a Reader of digital engineering with the School of Physics, Engineering and Technology, University of York, U.K., where he is the Director of iTWINS Lab and the Deputy Pillar Lead for Advanced Communications with the Institute for Safe Autonomy. He was a SINGA Ph.D. Scholar with the Institute for Infocomm Research, A-STAR, National University of Singapore. Then he worked at different academic and industrial positions in the Republic of Ireland and U.K. He has published a book and more than 100 peer reviewed book chapters, journal, and conference papers. He has been the Networks Working Group Chair of COST Actions CA15104 (IRACON) and CA20120 (INTERACT). He had chairing roles in organizing and technical programme committee of several IEEE flagship conferences, including IEEE ICC 2024, EUCNC 2025 and 2019, respectively, and PIMRC 2024 and 2019, respectively. He is also the Treasurer of the IEEE U.K. and Ireland Diversity, Equity, and Inclusion Committee. He is an Associate Editor-in-Chief of IEEE Communication Standards Magazine and a Fellow of U.K. Higher Education Academy. His current research interests include the application of machine learning in wireless networks, open radio access and networking, green and sustainable networks, airborne networks, digital twins of networks, and the Internet-of-Things.