

An intelligent platform: The use of O-RAN's SMO as the enabler for openness and innovation in the RAN domain

Content

Introduction	3
Automation in network operations	4
Network automation and Open RAN	6
SMO as a driver for RAN optimization	10
Ericsson recommendations	15
Glossary	17
References	19
Authors	20

Introduction

Telecom network operations are becoming increasingly complex, owing to increasingly diverse and demanding use cases, increased frequency band choices, virtualization, network slicing, as well as the ongoing disaggregation and distribution of the radio access network (RAN) functions, such as through Open RAN.

The Open RAN Alliance's (O-RAN) service management and orchestration (SMO) platform is an intelligent automation platform which applies automation at scale to simplify the complexity of networks, as well as improve network performance, enhance customer experience and minimize RAN operational costs.

In this white paper, we explore the new SMO and how it can benefit today's and tomorrow's RAN operations. The SMO enables automation and increases the abstraction offered to the users by managing RAN as a service and intents. We also provide our recommendations on the future specifications of SMO, as well as optimal deployment scenarios.

Automation in network operations

Telecom networks are becoming immensely powerful and generating unprecedented new revenue possibilities for communications service providers (CSPs). However, this is also bringing new levels of complexity to network operations, making automation a critical component for managing today's network operations.

In recent years, the exponential rise in mobile applications and services, such as media streaming, have leveraged the full capabilities of the network and, at the same time, generated huge volumes of data traffic. This has seen CSPs migrate from managing and delivering simple voice services to multiple data services, and now to offering complex consumer and enterprise services often with unique service-level agreements (SLAs). This proliferation of new and more complex services is one reason why automation has become critical in network operations.

Another reason for automation is the significant and complex transformation which has and is taking place across the network infrastructure. The technology evolution with 4G and 5G networks has enabled mass-market, complex data services to be delivered at scale and with high levels of reliability, however this has also led to multiple technology stacks and RAN generations within the networks. In addition to this, the widespread adoption of small cell based private and public networks has hugely increased the number of network elements that need to be managed, as has the increased number of frequency bands required to support higher maximum user throughput and increased capacity.

The evolution from physical networks, to virtualized, containerized and now fully cloud-native networks has been transformational. Many CSPs have networks comprising physical

equipment, virtualized functions and even major functions such as the 5G core deployed as cloud-native components in hyperscale cloud environments – resulting in a wide range of flexible deployment options. Increasingly, CSPs are also deploying capabilities such as network slicing and closed-loop assurance to expand their network’s capabilities and improve its quality, driving a greater need for network automation. This increases the level of abstraction offered to the network operators, moving from managing a set of resources to managing towards wanted objectives or intents.

A perfect storm – factors that drive the need for automation

<p>Manage the increasing network complexity created by multiple generations of radio access network (RAN) technology</p>	<p>Improve agility in managing and launching services across existing purpose-built and newer Cloud RAN/Open RAN infrastructures</p>	<p>Improve efficiencies of network operation by supplementing or replacing traditional manual network operation tasks</p>
<p>Overcome the skills gap and challenges pertaining to the manual management of new, complex cloud-native and physical networks</p>	<p>Secure new enterprise revenues requiring stringent SLA management and customization</p>	<p>Meet the ongoing rise in customer experience expectations to reduce churn and improve network promoter score (NPS)</p>




Figure 1. The automation "tipping point" has made network automation a critical part of network operations management

The GSMA's 2020 State of Mobile Internet Connectivity report [1] recently found that:

90 percent of CSPs believe that they must standardize and automate service activation and fulfilment processes, and 83 percent believe that they must support automated closed-loop service fulfilment and assurance.

Network automation and Open RAN

In coming years, the deployment of new 5G and eventually 6G use cases – together with the increasing disaggregation and distribution of the network's software and hardware components – will serve to intensify the current challenges related to operational expenditure (OPEX) efficiency and revenue growth.

Managing network operations accounts for largest share of OPEX across today's CSPs, as has been confirmed in recent studies by Analysys Mason [2] and Omdia [3]. Management of RAN operations is estimated to account for approximately 13 percent of overall network OPEX, however when combined with deployment, optimization, field maintenance and power expenses, overall RAN operations are estimated to account for almost 50 percent of the total network's OPEX.

As a result, CSPs are increasingly exploring ways in which to reduce RAN OPEX, such as the adoption of [open RAN](#) architecture and automation.

Based on internal analysis, together with customer discussions and experience from managed services, network operations OPEX split is as follows:

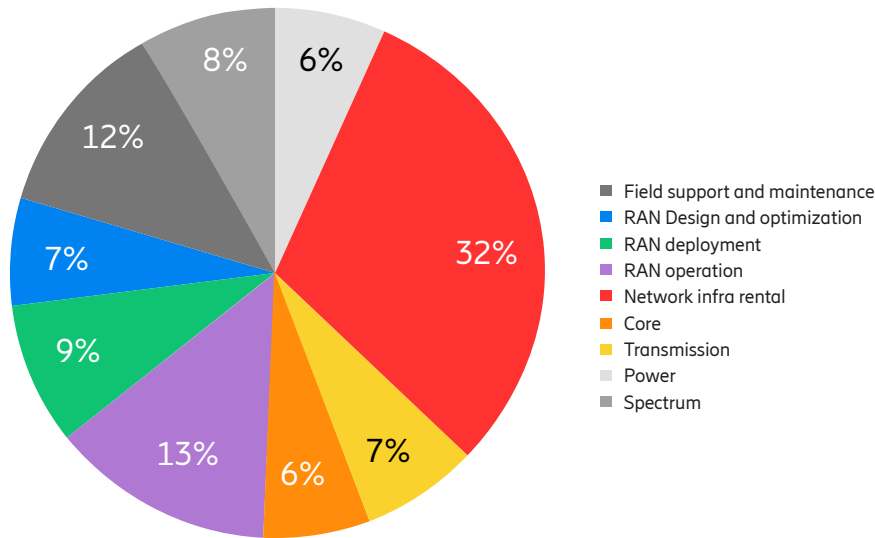
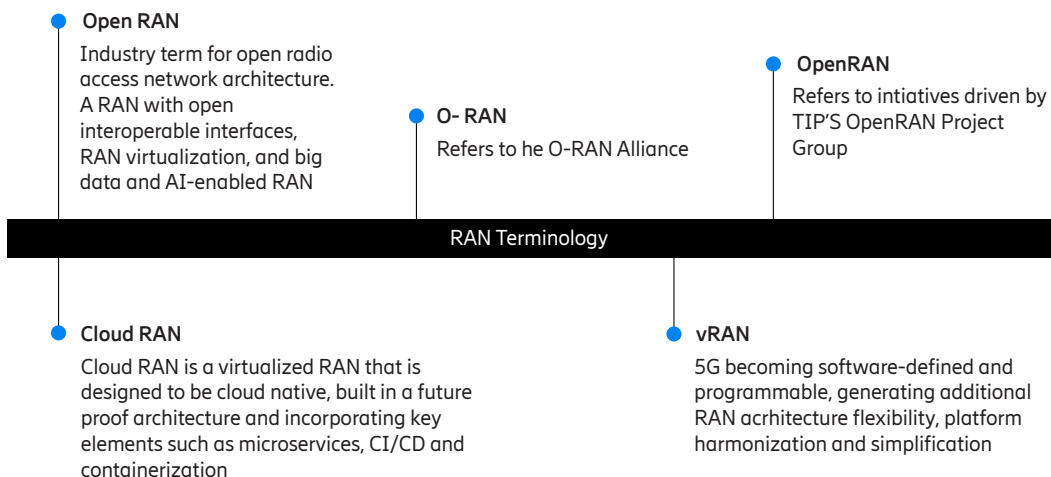


Figure 2. Typical share of OPEX per network domain and expense
 (Source: Ericsson internal)

Introduction to Open RAN

Current RAN technology is provided as a hardware and software integrated platform. The ambition for Open RAN is to create a multi-supplier RAN solution that allows for the separation - or disaggregation - between hardware and software with open interfaces and virtualization, hosting software that controls and updates networks in the cloud. The promised benefits include supply chain diversity, solution flexibility, and new capabilities leading to increased competition and further innovation.

There are multiple flavors of Open RAN technology, including Ericsson’s Cloud RAN.



Drivers of Open RAN

CSPs are pursuing Open RAN with the following underlying motivations and drivers:

- Openness and disaggregation to ensure supply chain security and reduce RAN capital expenditure (CAPEX)
- Cloudification for leveraging commercial off-the-shelf (COTS) cloud infrastructure to improve CAPEX efficiency and reduce RAN software (SW) operational and maintenance costs
- Accelerate innovation to drive new 5G revenues by opening up RAN control and management planes

Technology and industry changes in the RAN domain are increasing the complexity of network operations which impacts operational costs. Intelligent RAN automation, often using artificial intelligence (AI) and machine learning (ML) techniques, can enable CSPs to address traditional complexity challenges and optimize OPEX.

To enable this automation function, a new approach to network management and orchestration is required. Purpose-built RAN management, based on proprietary solutions and traditional vendor-specific approaches, will remain, however, will be enhanced by a new set of management and orchestration capabilities.

We believe that the SMO platform-based capabilities can be extended to handle multi-vendor RAN infrastructure across the CSP network. The platform will also automate the RAN by enabling applications from network vendors, CSPs and other third-party developers to improve network performance, enhance customer experience and reduce operational costs.

Ericsson analysis and proof-of-concept results:

SMO-driven automation can save brownfield CSPs up to 10 percent of their RAN OPEX over a period of five years.

Service management and orchestration in O-RAN

The Open RAN Alliance (O-RAN) defines technical specifications and interfaces related to the RAN's service management and orchestration (SMO) framework.

The SMO platform is an automation platform for Open RAN radio resources. Hierarchically, it is a component of the operational support system (OSS). Within the zero touch service management European Telecommunication Standards Institute (ETSI-ZSM) it is viewed as a RAN domain controller. The platform can be distributed, being deployed on premises, on the (telco) cloud or as-a-Service to suit the end-user requirements.

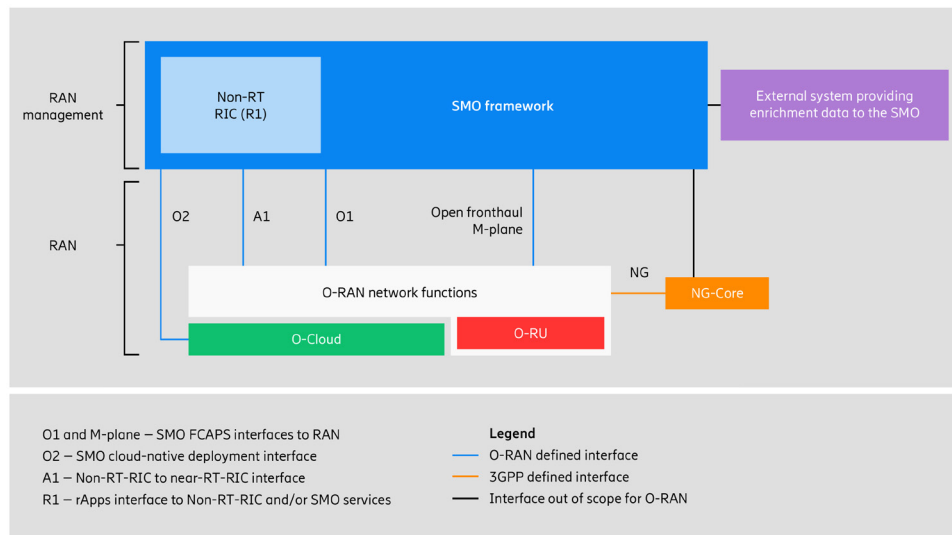


Figure 3. Service management and orchestration (SMO) platform

As defined in O-RAN's SMO framework, the network's open central unit (O-CU) functions, open distributed unit (O-DU) functions, and near-real-time radio intelligent controller (Near-RT-RIC) are defined as cloud-native virtualized functions which run on cloud infrastructure also known as the O-Cloud.

O-RAN is defining means of extending the O1, A1 and R1 interfaces to enable a competitive ecosystem and quick time to market of new functionality.

The architecture defines various SMO interfaces, namely O1, O2, and A1, which allow the SMO to manage multi-vendor Open RAN networks order support innovation as well as openness, ORAN is standardizing means of extending the O1, A1 and R1 interfaces to enable a competitive ecosystem and quick time to market of new functionality.

The O-RAN defined SMO includes:

- A design environment for rapid application development
- A common data collection platform for management of RAN data as well as mediation for O1, O2 and A1 interfaces
- Support for licensing, access control and AI/ML lifecycle management, together with legacy north-bound interfaces (NBI)
- Existing OSS functions such as service orchestration, inventory/topology and policy control

R1 interface to allow portability and lifecycle management of rApps. We believe that by supporting specific proprietary south-bound interfaces (SBI) to third-party equipment management systems (EMS) the SMO would have the ability to automate existing, purpose-built multi-vendor RAN as well as Open RAN networks.

A full description of the O-RAN architecture is available at www.o-ran.org/specifications.

SMO as a driver for RAN optimization

SMO is designed to be open by creating an automation platform, based on cloud-native principles, that deploys RAN functions and applications over open interfaces.

The O-RAN defines interfaces to enable inter-operability between the SMO, RAN functional entities and applications.

The Non-RT RIC function supports intelligent RAN optimization in non-real-time by providing policy-based guidance using data analytics and AI/ML training/inference. Non-real time is defined as an automation loop greater than one second. The non-RT RIC can leverage SMO solutions such as data collection and provisioning services of O-RAN nodes. rApps are modular applications that leverage the functionality exposed by the non-RT RIC/SMO Framework over the R1 interface to perform multi-vendor RAN optimization and assurance.

One of the expected benefits of this open architecture is to bring new application developers into an expected application eco-system. rApps can be developed and delivered by any third party. We expect rApps to be built by CSPs directly, but also that some rApps will be specified by CSPs and built by third parties, as well as some developers creating unsolicited new applications. Because rApps are based on open interfaces and are application platform agnostic, it means they can run on any vendor's non-RT-RIC.

The Non-RT RIC can access data from different network domains such as RAN, core, transport, as well as other external data sources. The non-RT RIC can also use data provided or enriched by rApps themselves. This makes the correlations and decisions done at non-RT RIC much more accurate with a broader visibility and insights to the network performance.

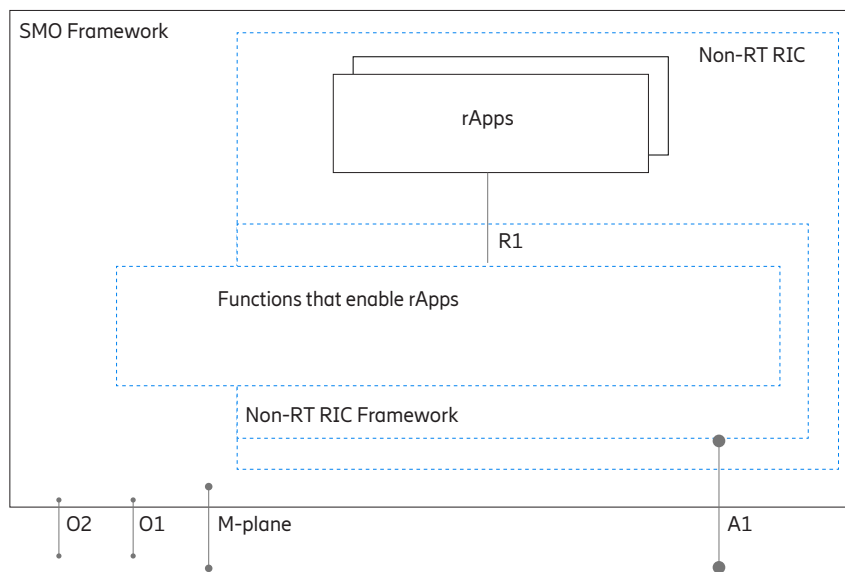


Figure 4. Logical interfaces of SMO

SMO interfaces explained

R1 interface operations towards multi-vendor rApps.

R1 interface is motivated by the need for portability of multi-vendor rApps and provides value-added services to rApp developers and solution providers. It is a collection of services including, but not limited to, service registration and discovery services, authentication and authorization services, AI/ML workflow services, and A1, O1 and O2 related services. The interface enables new Open APIs to be integrated in the SMO framework.

A1 interface for policy guidance

SMO provides fine-grained policy guidance such as getting User-Equipment to change frequency, and other data enrichments to RAN functions over the A1 interface.

O1 interface for operations and maintenance

SMO supports O1 interface for managing the operation and maintenance (OAM) of multi-vendor Open RAN functions including fault, configuration, accounting, performance and security management, software management, and file management capabilities.

O2 interface for cloud infrastructure management

SMO supports O2 interface to support cloud infrastructure management and deployment operations with O-Cloud infrastructure that hosts the Open RAN functions in the network. O2 interface supports orchestration of O-Cloud infrastructure resource management (e.g., inventory, monitoring, provisioning, software management and lifecycle management) and deployment of the Open RAN network functions, providing logical services for managing the lifecycle of deployments that use cloud resources.

M-plane

For supporting multi-vendor O-RU integrations, SMO supports NETCONF/YANG based Open FrontHaul M-Plane as an alternative to O1 interface. Open FrontHaul M-plane supports the management features including startup installation, software management, configuration management, performance management, fault management and file management.

Extending the O-RAN definition of SMO

While there is significant value in aligning on an open industry standard full maturity and compliance will take time. Therefore, some vendor extensions of the O-RAN interfaces will have to be supported to deliver overall automation value in the RAN. Model-driven approach in SMO is an efficient way in handling multi-vendor integrations. With this approach, unified models, vendor specific extensions and schema changes are supported without any SMO software changes. This removes the re-integration cost between the SMO and multi-vendor RAN functions.

As the SMO concept in O-RAN allows for improved closed loops compared to today's multi-vendor centralized self-organizing network (SON) solutions, as well as the possibility to integrate with other OSS and business support systems (BSS), we believe new innovative automation use cases will be coming from the SMO application ecosystem community. We also see a possibility of migrating market leading proven solutions such as AI-based multi-vendor network optimization modules to an SMO as rApps.

The O-RAN defined SMO entity can be extended to a broader radio automation platform including purpose-build RAN management through the A1 interface. We believe SMO plays an active role in realizing the end-to-end use cases by supporting multi-domain orchestrators to provide a single solution for CSPs' nationwide multi-vendor RAN networks. For example, an SMO, together with one or several rApps, may need to secure that latency is kept within allowed limits for a mission-critical network slice.

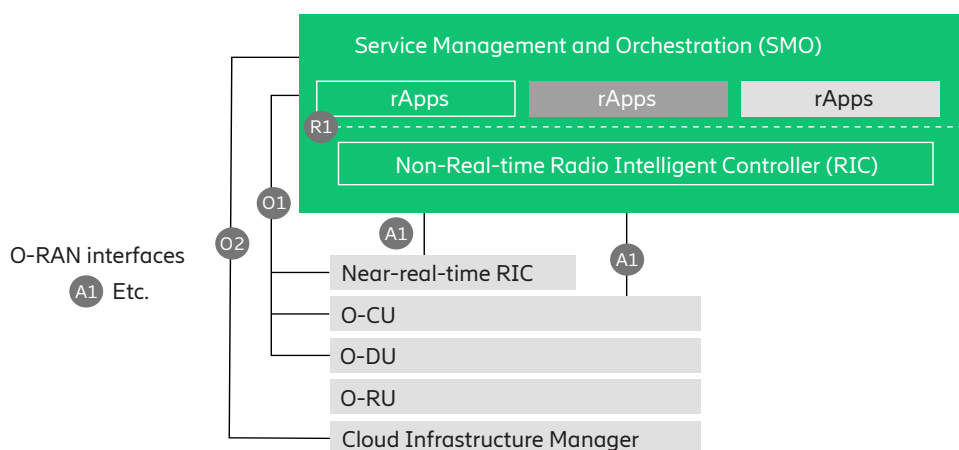


Figure 5. O-RAN Alliance architecture with SMO and key interfaces

Multi-vendor RAN support

The role of SMO in standardized development and execution environments is to support multi-vendor RAN.

SMO is an automation platform for Cloud RAN/Open RAN and purpose-built RAN. SMO offers platform capabilities using which automation applications can be built to support most common management use cases for multi-vendor RAN networks. The application development and run-time environment in SMO offers a structured and standardized way to address these automation-led network operations:

- The design environment enables rapid application development and automation artifacts.
- The run time environment enables multi-vendor applications to run on the platform by offering capabilities such as Application life-cycle management, AI/ML training and execution, Security, Data management, Policy, Analytics, Service and resource Orchestration etc. The capabilities are exposed to the multi-vendor applications (rApps) over the R1 interface.

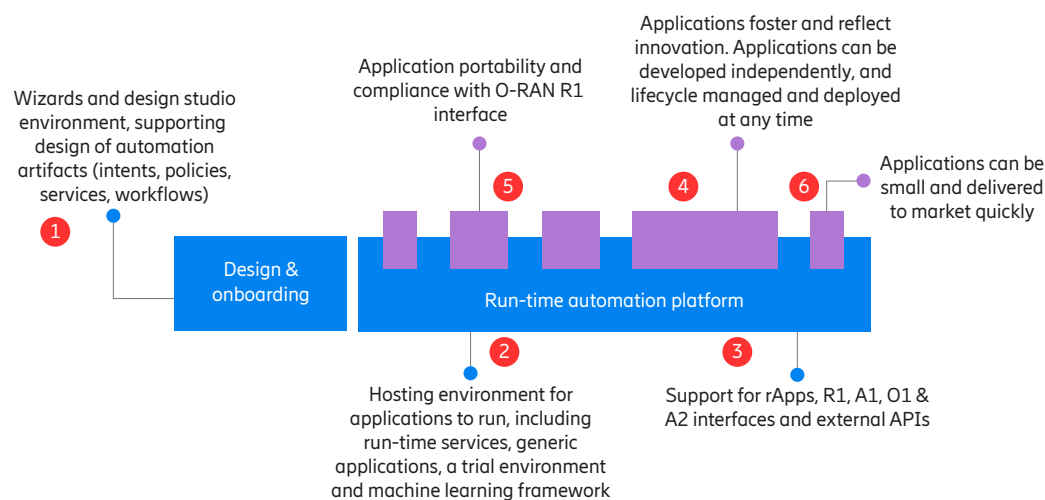


Figure 6. SMO application development and run time environments

The capabilities of the underlying platform are exposed to the multi-vendor rApps over the R1 interface. The platform capabilities exposed on the R1 interface include service exposure, security, data management, automation, analytics insights, ML execution and training, service and resource orchestration, topology and inventory, application lifecycle management, controller framework etc.

Application lifecycle management

To support multi-vendor O-RAN applications, the SMO provides a unified approach towards common lifecycle management. This is applicable for all cloud-native applications (rApps, xApps) and cloud-native network functions (O-CU-CP, O-CU-UP, O-DU, O-RU etc.).

Security

As networks become pervasive, more dynamic and more intrinsic to daily life, security becomes ever more critical. The SMO, as a gateway to the network, is a critical function from a security perspective, and security in all aspects of the software development, delivery, deployment, and optimization flows is critical, both for the networks managed by the SMO as for the SMO itself.

Data capabilities

Data is becoming ever more important in automation solutions. Developers want timely access to relevant data, and the SMO is a key enabler for providing these capabilities towards the applications. This includes data collection, management, and movement capabilities.

SMO targets "data on demand" capabilities for the development of automation features through well-defined interfaces

- Data location discovery for spot data access
- Data request and subscription for real time and batch delivery, filtered by data sources, type, geo location, etc.
- Target and role-based access control to data, regardless of the delivery mechanism
- IT mainstream cloud native interfaces for real time and batch data

The same capabilities can be used by automation features to supply data and insights to SMO and expose them back to other features.

Automation

One of the key value propositions of O-RAN Alliance architecture is enabling intelligent automation towards the RAN using multi-vendor applications. The non-RT RAN Intelligent Controller is a logical function in the SMO framework that enables non-real-time control and optimization of RAN elements and resources. The SMO offers a wide range of automation capabilities as part of the platform and enables portable and smaller applications to deliver value-added functionality.

Service and resource orchestration

The service and resource orchestration function in SMO orchestrates the placement and deployment of containerized network functions and applications on the O-Cloud based on the required characteristics of the service to be deployed. An efficient O-Cloud cluster selection and lifecycle managing the cloud resources is key to achieving improved performance.

Mediation capabilities

SMO supports mediation of RAN functions via O-RAN interfaces that include but not limited to modification of configurations, collection of data related to configuration, fault management, performance management etc.

Ericsson recommendations

While many CSPs are currently running Cloud RAN/Open RAN trials and proof-of-concepts in rural areas as standalone greenfield networks, we believe such deployments do not offer an optimal representation of the full benefits of SMO and its impact on network performance, customer experience and operational costs.

The key role of the SMO is to apply automation at scale to simplify the complexity of disaggregated Open RAN networks.

If the capabilities of SMO can be extended across multi-vendor and, more importantly, multi-technology networks including networks defined as Cloud RAN/Open RAN and purpose-built 4G and 5G networks, it would offer a powerful platform to significantly transform automation of mobile networks.

We make the following recommendations for SMO:

1. Deploy in complex RAN environments - Our recommendation therefore is to deploy SMO into complex multi-vendor, multi-technology networks where the automation platform can create the greatest operational impact.
2. Target addressable areas of high operational costs – Owing to network deployment, field support and maintenance (12 percent network OPEX), as well as network operation, RAN optimization (13 percent network OPEX), representing two of the largest addressable areas of OPEX, we recommend that network deployment and network optimizations are early targets for new automation rApps and xApps.
3. Leverage proven use cases – The ability to take existing centralized SON and network design and optimization applications and use them to create new automation applications running on a common platform is seen as an opportunity to take proven capabilities and deploy them at scale.

4. Use best in class applications – Use automation applications from a wide selection of developers including network equipment vendors, CSPs and third-party specialists. Be prepared to test similar applications and select those that provide the best ROI. Utilize anything the platform vendors can do to accelerate open innovation such as software development toolkits, partner development ecosystems and even marketplaces to maximize choice.
5. Deploy across multi-technology networks – SMO provides maximum return on investment (ROI) where it enables automation for 100 percent of the network. Having high-degrees of automation in an Open RAN/Cloud RAN network is a way of ensuring the new technology is future-proof. However, adding high degrees of automation to the existing purpose-built RAN that makes up more than 95 percent of CSP networks today and more than 99 percent of cost is also highly desirable.

The value of SMO is in automating and orchestrating RAN complexity, irrespective of multi-vendor or multi-technology networks. Ultimately, the SMO should provide a single pane of glass for RAN automation as well as enabling open rApp innovation with tools such as an SDK will genuinely reduce complexity in these evolving networks.

The aspiration for all communication service providers is an SMO entity which fits seamlessly into their existing architecture and supports network slicing, assurance and complex service orchestration.

Glossary

A1: Interface between non-RT RIC and near-RT RIC to enable policy-driven guidance of near-RT RIC applications/functions and support AI/ML workflow.

E2: Interface connecting the near-RT RIC and one or more O-CU-CPs, one or more O-CU-UPs, and one or more O-DUs.

Near-RT RIC (Near-real-time RAN Intelligent Controller): A logical function that enables near-real-time control and optimization of RAN elements and resources via fine-grained (for example, UE basis, cell basis) data collection and actions over E2 interface.

Non-RT RIC (Non-real-time RAN Intelligent Controller): A logical function that enables non-real-time control and optimization of RAN elements and resources, AI/ML workflow including model training and updates, and policy-based guidance of applications/features in near-RT RIC.

O1: Interface between orchestration and management entities (orchestration/NMS) and O-RAN managed elements for operation and management, by which FCAPS (fault, configuration, accounting, performance and security) management, software management, file management and other similar functions shall be achieved.

O2: Interface towards the O-Cloud to support orchestration of O-Cloud infrastructure management and deployment of the network functions on the O-Cloud.

O-Cloud: O-RAN Cloud computing platform comprising a collection of physical infrastructure nodes that meet requirements to host the relevant O-RAN functions (such as near-RT RIC, O-CU-CP, O-CU-UP and O-DU, and so on), the supporting software components (such as operating system, virtual machine monitor, container runtime, and so on) and the appropriate management and orchestration functions.

O-CU-CP (Central unit – control plane): A logical node hosting the RRC and the control-plane part of the PDCP protocol.

O-CU-UP (Central unit – user plane): A logical node hosting the user-plane part of the PDCP protocol and the SDAP protocol.

O-DU (Distributed unit): A logical node hosting RLC/MAC/High-PHY layers based on a lower layer functional split.

R1: Interface comprised of a collection of services that facilitate the interactions between rApps and the non-RT RIC framework.

rApp: An application designed to run on the non-RT RIC. Such a modular application leverages the functionality exposed by the non-RT RIC to provide added value services relative to intelligent RAN optimization and operation. rApps communicate with the platform via the R1 interface.

References

1. ["The State of Mobile Internet Connectivity 2020"](#). GSMA. 2020
2. ["Open RAN: ready for prime time?"](#). Analysys Mason. April 2021
3. ["OSS/BSS Evolution Survey – 2022"](#). Omdia. November 2020
4. ["Open RAN: weighing the challenges versus the promise"](#). GSMA Intelligence. June 2021

Authors



Roger Niklasson is an experienced Strategic Product Manager working in Solution Area OSS as Ericsson. He has more than 20 years of experience in the Telecom domain whereof the last 10 years within OSS.



Rjavarma Bhyrraju is a Portfolio Architect for Standards and Open-Source models in Solution Area OSS in Ericsson, where he works in transformation of OSS products with emerging technologies and trends in the software, telecommunications and network management. He has more than 22 years of experience in standardization and product development.



Kedar Thakar is Principal Consultant in the Strategy and Business development organization of Ericsson Digital Service Business Area, where he is responsible for evaluating strategic options for OSS offerings in terms of new market entry, go-to-market strategy, commercial strategy, partnerships and M&A. He has over 25 years of telecom experience.



David Espadas is responsible for Product Marketing in Solution Area OSS at Ericsson, where he works to build the best market perception for the comprehensive Ericsson OSS portfolio covering Network Management, Orchestration and Analytics capabilities. He previously spent more than 10 years in business development and presales in those same areas plus Business Support Systems (BSS).



Gustavo Hylander is a Strategic Product Manager with more than 10 years of experience. Currently responsible for the Centralized SON products in Solution Area OSS in Ericsson. With more than 20 years of experience in the Telecom domain. He has mainly worked in the network design and optimization domain.



Justin Paul joined Ericsson Digital Services as Senior Solutions Marketing Manager OSS in 2020. Since joining Ericsson, he has been involved in the O-RAN Alliance Service Management and Orchestration domain.