

Open Radio Access Networks (O-RAN)



Partha Seetala
Robin.io

Rapid scaling of profitable 5G Services with Cloud-Native orchestration

Partha Seetala, Founder, Chief Executive Officer, Robin.io

Partha Seetala, Founder and CEO, Robin.io

Partha Seetala is the Founder and Chief Executive Officer at Robin.io. With more than 20 years of technology and product vision and expertise, he has taken multiple fully-featured products from concept to market. He has conceived, designed and built products in Scale-out distributed storage, File Systems, Networking, Distributed Systems, Big Data, and Information Analytics and Containers/Kubernetes space.

Previously he was a Distinguished Engineer at Symantec's information and storage management business, known as Veritas. In that capacity, he conceived, architected and led engineering teams to take multiple products from concept to market. He was also an adviser on multimillion-dollar product lines including NetBackup, Cluster File System, Veritas Cluster Server, and Information Fabric. He holds a master's degree in Computer Science and Engineering from the University of Minnesota.

Abstract:

Since 5G service requirements span so many industries, they are both a business opportunity and challenge, requiring one to deploy innovative mobile services faster than ever before. Unlike the preceding three pillars, which all add complexity to the equation when appropriately executed, orchestration and automation can make it all appear effortless to the end-user and the operator.

Your business agility will ultimately rely on your orchestration platform's flexibility, ease of use, and breadth of scope. To keep up with 5G's complexity, the design of your network, services, and how they work together needs to become more dynamic and customizable across the entire lifecycle, keeping up with architectural changes and market demands. This operations mandate dictates reduced time to outcome across multiple fronts, including continuous development, testing, and deployment.



Summary

The countdown is on to shift from legacy solutions to modernized, agile and efficient networks with massive scale that enables a wealth of 5G services and applications. Operators and providers worldwide are gaining momentum in their 5G deployments by choosing fully orchestrated, cloud-native platforms. According to GSMA, by 2025, there will be 1.2 billion 5G connections worldwide.

Flexible and adaptive cloud-native platforms and orchestration tools will help operators explore a more competitive and vibrant supplier ecosystem with faster innovation. This will lead to deeper market penetration and improved user experience. By utilizing advanced Kubernetes platforms, with unified operations models and fully shared resource pools, the limitless potential of 5G is possible, creating new business revenue while saving costs through advanced lifecycle automation.

Change is leading the way with four paradigm shifts powering innovative and exciting new 5G services.

Edge Data Center Proliferation

The spectrum of new 5G services requires significantly more bandwidth, delivered at much lower latency. This drives a massive proliferation in the number of cell sites and edge data centers. These represent significant investments in infrastructure and operations platforms. 5G data center proliferation increases the number of cell towers by over 100x and user density by almost 1,000x. Even simple math dictates an explosion of integration touchpoints, as well as the number of lifecycle operations completed every minute. Furthermore, edge data center capacity will be at a premium and spread across many locations that by yesterday's standards are considered remote. What used to be a walk down the row of a data center will now become a truck roll unless the entire solution stack, bare-metal to services, can be efficiently orchestrated at scale.

Open Radio Access Networks (O-RAN). The O-RAN Alliance's mission is to "re-shape the RAN industry towards more intelligent, open, virtualized and fully interoperable mobile networks." O-RAN marks the transition from proprietary vendor lock-ins to open models that enable operators to select RAN hardware and software from multiple vendors, enabling them to be more adaptive over a wide array of markets, customers and solutions. A diverse and vibrant RAN supplier ecosystem with faster innovation will facilitate accelerated market penetration and improved user experience.

While the benefits of O-RAN to new services are enormous, they come at the cost of solution integration tasks in terms of interoperability and Network Function (NF) performance tuning. It is now more important than ever to rapidly automate the innovation and deployment pipelines.

Multi-access Edge Compute (MEC)

Coupled tightly with O-RAN is MEC, an interoperable framework for distributed edge computing. A distinguishing feature of MEC applications, besides providing low-latency services that can be tailored regionally, is that they can interactively contribute to O-RAN optimization, delivering faster, localized decisions that improve both data transport performance and service experience. Leading MEC applications include content delivery, interactive real-time gaming, connected transportation, smart cities and smart venues. In many cases, these MEC applications will be deployed alongside and even in the same cluster as O-RAN NFs. This opens additional revenue generation opportunities via edge application hosting. Upfront planning will be essential as this adds other complexities in security, multi-tenant isolation, autoscaling, monitoring, role-based access, and granular chargeback.

Orchestration and Automation



Since 5G service requirements span so many industries, they are both a business opportunity and challenge, requiring one to deploy innovative mobile services faster than ever before. Unlike the preceding three pillars, which all add complexity to the equation when appropriately executed, orchestration and automation can make it all appear effortless to the end-user and the operator.

Your business agility will ultimately rely on your orchestration platform's flexibility, ease of use, and breadth of scope. To keep up with 5G's complexity, the design of your network, services, and how they work together needs to become more dynamic and customizable across the entire lifecycle, keeping up with architectural changes and market demands. This operations mandate dictates reduced time to outcome across multiple fronts, including continuous development, testing, and deployment.

As operators are in the midst of their cloud-native transition, multi-vendor solutions come hand-in-hand with multiple deployment models, namely containers and Virtual Machines (VMs). Your cloud platform needs to support both. Still, it needs to be done in a way that maximizes performance for all models while truly eliminating not just resource siloes, automation silos and operations silos. During your automation planning, one must be rigorous as it is easy to create new operations silos to remove others.

This begs the question: how do you choose the right cloud platforms and orchestration tools that best deliver your 5G solutions, avoiding potential pitfalls? The answer is, you dig deeper into the solution and understand HOW, not just what. Here are some key questions to ask:

- How do I unify my bare-metal with my services life cycle? Can I streamline these life cycles with a single tool with integrated bare-metal, cluster, and services workflows?
- How easy is it to remotely tune bare-metal for performance, pinning resources based on CPU, GPU, memory, HugePages, underlay/ overlay networks and affinity/ anti-affinity?
- What granularity and breadth does my automated workload placement algorithm control? Can it be configured to act on the node, pod, cluster, data center and multi-cloud levels?
- Can I easily or even automatically retune application policies for caveats in configuration, resources, and environment, as I don't want to turn my automation platform into a manually intensive scripting nightmare?
- How many steps does it take, if any, to reuse workflows and resource pools when deploying VMs and containers side-by-side?
- How can I give my VMs the same operational and performance benefits that my containers have?
- As I add functionality, am I just piling on more licenses, software packages, and operational caveats to an old system that was never really designed for containers in the first place?
- When I migrate, clone, snapshot, backup, and restore, how easy is it? What do I get – do I have a copy of the entire application, database, and the relationships between all of the components that make it up? When I clone my service, is there pre- and post-processing to ensure my clone automatically deploys as a working copy? Or do I have to perform additional manual tasks, so it doesn't think it is the original?
- How declarative is my platform? Do I need to be a Kubernetes expert, or does the interface know my service's NF composition and other supporting applications that have to be pre-installed before my service, databases, message queues, and the like? How easy is it to onboard all of this intelligence?
- How are resources reserved for multi-tenant deployments? What views and tools are available on a per-tenant basis?
- How customizable are your roles-based access definitions?
- How easy and granular are my chargeback options, and what is the scope?
- How does it help me with resource and outage planning across multiple events?
- What is the granularity and breadth of the monitoring framework, and how many extra vendors do I need to add to make it work the way I want it to?



Conclusion

When moving from the test lab to deployment and massive scale-out of 5G services, it will eventually become evident that HOW you automate is just as important as what you automate. There is a lot of hand-waving in the industry where vendors and analysts praise the benefits of the cloud automation revolution as if it is a simple cure-all for any repetitive or scale-out task. Automated cloud platforms and orchestration solutions are in and of themselves, their own deployments with significant comparative benefits. Choosing these platforms must be done with great care as they impact your time to the outcome, resource utilization, solution costs, and opportunities.