



Cloud Native for Telco: Making IT Technology Feasible at the Network Edge

WHEN IT MATTERS, IT RUNS ON WIND RIVER

EXECUTIVE SUMMARY

The term *cloud native* has come to dominate telco industry discussions around network transformation. Communications service providers (CSPs) are realizing that the cloud-based networking practices of the world’s leading web-scale companies can help their efforts become more efficient and deliver compelling end-user services.

Since Network Functions Virtualization (NFV) was introduced in 2012, many CSPs have pursued rapidly evolving network virtualization strategies: Physical network functions (PNFs) are being replaced by virtual network functions (VNFs) running in virtual machines (VMs) on industry standard IT hardware, rather than legacy appliances. Nonetheless, CSPs recognize that traditional virtualization is not enough to deliver significant cost savings and meaningful operational efficiency. They are focusing on cloudification and the implementation of cloud native services designed specifically for cloud environments.

CSPs are simultaneously pursuing edge computing strategies, distributing intelligence and processing resources to the edge to help deliver new experiences for customers and meet requirements for 5G services. Distributed edge cloud deployments require more operational flexibility and simplicity than centralized data centers, which makes them ideally suited for cloud native services.

But just because the term cloud native is pervasive doesn’t mean everyone has a firm grasp of the concept. And even if CSPs have already started down the path to virtualization by deploying VMs, they don’t have to start over to adopt cloud native technologies. This white paper aims to clarify what cloud native means, why it is necessary for edge cloud deployments, and how it presents solutions for implementing a cloud native strategy, wherever CSPs are in their network transformation.

TABLE OF CONTENTS

Executive Summary2

What Is Cloud Native?3

Benefits of Cloud Native Environments4

Challenges of Edge Cloud Deployments5

How to Migrate to Cloud Native6

How to Build Your Own Cloud Native Edge Cloud with StarlingX7

Conclusion7



WHAT IS CLOUD NATIVE?

The best-articulated and industry-accepted definition of cloud native comes from the Cloud Native Computing Foundation (CNCF), an open source project working on making cloud native computing sustainable. The CNCF defines the term as follows:

Cloud native computing uses an open source software stack to deploy applications as microservices, packaging each part into its own container, and dynamically orchestrating those containers to optimize resource utilization.

The CNCF also explains why cloud native components are so important for cloud computing: “Cloud native technologies empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid clouds.... These techniques enable loosely coupled systems that are resilient, manageable, and observable. Combined with robust automation, they allow engineers to make high-impact changes frequently and predictably with minimal toil.”

In other words, cloud native technologies enable telcos to run applications that are highly scalable, resilient, and easily managed via automated operations in any cloud environment.

Cloud native can be misconstrued as being synonymous with containers. While containers are essential to cloud native environments, they are not the only component of a cloud native solution. Cloud native technologies include containers, service meshes, microservices, immutable infrastructure, and declarative APIs that allow deployment in cloud environments through loosely coupled and automated systems. Containers offer a fast, lightweight, and disposable method for deploying and updating workloads and devices at the edge, and they can be used independently of a public or private cloud.

Today, containers are predominantly deployed by the telco/IT sector, according to Allied Market Research, and the sector is likely to be by far the largest adopter of containers by 2025. Over the next five years, other industry verticals are also expected to increase their reliance on containers, including banking and financial services, retail, healthcare, government, and education.

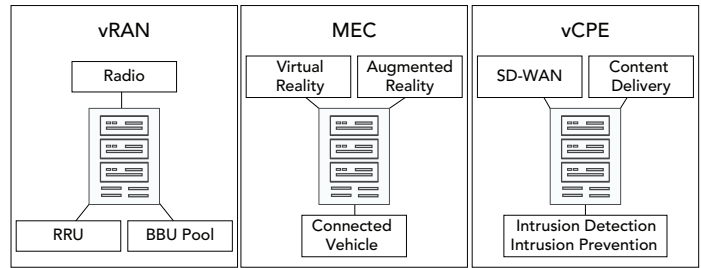


Figure 1. Telco edge use cases for cloud native environments

Edge Use Cases Need Cloud Native Environments

For the telco sector, the top use cases for containers and, more generally, cloud native techniques are:

- **Virtualized radio access network (vRAN):** Creates more flexible and efficient mobile networks by centralizing virtualized signal processing functions and distributing radio units at cell sites
- **Multi-access edge computing (MEC):** Distributes compute resources into access networks to improve service experience and support new latency-sensitive applications
- **Virtual customer premises equipment (vCPE):** A virtualized platform that delivers managed services to enterprises at lower cost and higher flexibility than hardware appliances

Some CSPs have already started to deploy these use cases in the cloud, mostly via VMs, and are looking to migrate to containers. Not surprisingly, many of the primary use cases for containers are also foundational to 5G, which generally will require a greater reliance on software-based networks and cloud native services.

But the use cases for 5G and cloud native environments are not limited to the telco sector. Cloud native network transformations offer CSPs the opportunity to expand into vertical markets where they have not previously played a significant role beyond providing basic connectivity services. CSPs can add more value and serve a wider variety of new customers through cloud native edge deployments.

In transportation, for example, there are compelling 5G use cases for smart transit and rail management systems, autonomous shipping, and autonomous vehicles. In the realm of Industry 4.0, companies are exploring 5G-enabled robotics applications, human-machine interfaces (HMIs), and virtual programmable logic controllers (vPLCs). The healthcare sector is also actively pursuing how to leverage 5G for imaging, monitoring, and diagnostics.

The use cases share some commonalities: They all need local compute and processing resources at the edge of the network; they must support high-bandwidth and/or low-latency applications; and they need workload consolidation. Also, they are critical infrastructure workloads, which means they have additional requirements for performance, security, and availability.

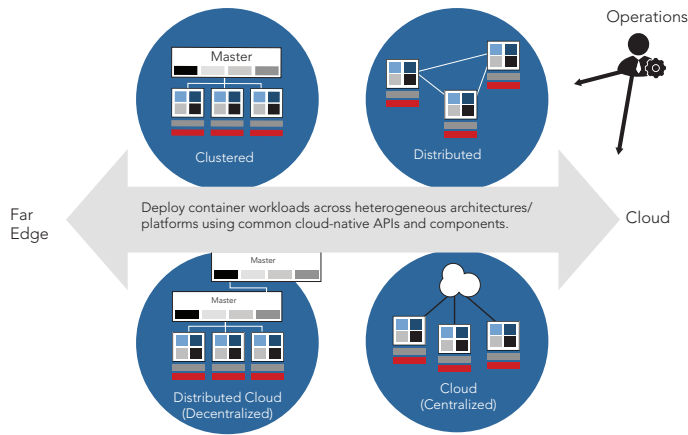


Figure 2. Heterogeneous edge environment

Navigating Heterogeneous Environments at the Edge

There are many different network edge locations within what could be called the “broad” edge: from the near edge in telco networks where base stations are deployed to the far edge on premise of factories, enterprises, or stadiums. As the edge evolves, the definition will likely evolve to include the devices themselves.

Across the broad edge network, there will be different cloud architectures deployed, depending on the use case and what legacy infrastructure is in place. For a campus or enterprise, the architecture could be a decentralized, distributed cloud where each site is autonomous. Other edge use cases may require a distributed hierarchy architecture or a centralized cloud where control functions are centrally located.

The result is a heterogeneous architecture environment at the edge, which is significantly different from the more homogeneous environment of a centralized data center.

Ultimately, cloud native technology helps CSPs navigate these heterogeneous environments by giving them “landing zone” autonomy at the edge. That is, regardless of the architecture type or edge location, a cloud native environment allows CSPs to deploy application container workloads across heterogeneous architectures and platforms using common cloud native APIs and components. From a single pane of glass, containers can be deployed to any architecture, any platform, and any location across the network edge.

BENEFITS OF CLOUD NATIVE ENVIRONMENTS

Building on the landing zone autonomy advantage, the biggest benefits of containers and cloud native environments are in the areas of deployment, operations, management, and development of applications. Altogether, cloud native fosters a DevOps approach to delivering applications and services.

More specifically, cloud native benefits include extra security and isolation from the host environment, easy version control in software configuration management, a consistent environment from development to production, and a smaller footprint. In addition, cloud native allows developers to use their preferred toolsets and enables rapid spin up and spin down as well as easy image updates. It facilitates portability to run on all major distributions while also eliminating environmental inconsistencies.

Depending on the role of persons involved in application deployment, cloud native presents different advantages, as follows:

- **Infrastructure provider:** Containers offer a logical packaging mechanism in which applications can be abstracted away from the environment, so that people involved in infrastructure provisioning need not be concerned with the specific workloads.
- **Application manager:** Cloud native allows quick, easy, and consistent application deployments, regardless of the target architecture, and replaces complex upgrades with ephemeral updates. Day 2 operations managers can easily handle the deployment, monitoring, and lifecycle of the workloads.
- **Developers:** Containers facilitate DevOps by easing the development, testing, deployment, and overall management of applications. Developers can design and test once and then deploy everywhere. Also, containers are lighter weight than VMs. They start up much faster and use much less memory compared to booting an entire operating system (OS).

CHALLENGES OF EDGE CLOUD DEPLOYMENTS

As CSPs consider edge deployment strategies, there are four main challenges to overcome:

- **Complexity:** With multiple use cases and a heterogeneous environment, edge cloud deployments are inevitably complex, especially in the areas of operations and management. Not all service providers have the resources to keep up with the fast-changing cloud native landscape, where new open source projects are launched almost weekly.
- **Diversity:** There is no one-size-fits-all approach for edge cloud deployments. There are very real differences between operational technology (OT) and information technology (IT) environments. For example, each has different latencies, bandwidth, and scale requirements, among other things. Service providers will need to bridge the gap.
- **Cost:** With potentially thousands of edge servers deployed, the hardware costs can exceed five times what it costs to deploy an application in a data center. In some cases, a worker node might require four servers to host control planes, storage, and compute, and these will require extra fans, discs, and cooling. The costs can easily become unfeasible.
- **Security and performance:** Since many edge sites don't have the physical security that a large data center would, security needs to be built into the application and infrastructure. And the edge clouds must meet telco and other critical infrastructure performance requirements.

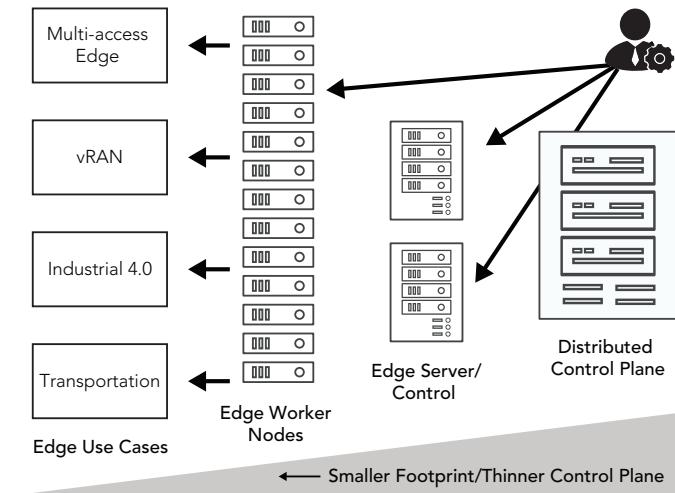


Figure 3. Distributed edge cloud topology

Before discussing the key requirements for edge cloud deployments that will overcome these challenges, it's important to establish what an edge cloud topology looks like. At the edge, there is likely to be a multitier, or n-tier, architecture with a distributed control plane, where the provisioning, lifecycle management, logging, and security are performed. There are also edge servers and worker nodes. The edge servers are like mini data centers with lightweight control planes that manage the worker nodes, which could be located at the base station in a mobile access network or on premise at an enterprise. The worker nodes run the application workloads, which can be anything from telco apps to OT-based workloads, as in factory automation or transit system management.

Where the workload is deployed depends on the latency requirements for the application. The lower the latency required (for example, from 5 to 20 milliseconds between device and worker node), the worker node will be deployed as far out to the edge as possible. As nodes are deployed further out, the requirement for lightweight control plane and smaller physical footprint becomes more important.

Given the topology of distributed edge clouds, there are certain requirements that CSPs must have to deploy edge cloud costs efficiently:

- **Workload orchestration:** CSPs need a clear determination of where and how an application workload should be processed across an n-tier gradient of compute, storage, and network resources provided by the edge cloud.
- **Zero touch provisioning:** An automated system is necessary for managing workload orchestration that will take into account the performance, time and cost requirements of an application.
- **Centralized management:** CSPs need a system-wide view of all the servers and devices in the edge network from a single pane of glass. Also, a centralized control plane can provide services such as logging, storage, security, updates, and upgrades as well as lifecycle management.
- **Edge cloud autonomy:** If connectivity is lost between the central site and an edge cloud site, the edge cloud needs enough control plane functionality in order to carry on performing mission-critical operations.
- **Massive scale:** The solution needs to be able to scale from tens to hundreds of thousands of edge clouds in geographically dispersed locations.

HOW TO MIGRATE TO CLOUD NATIVE

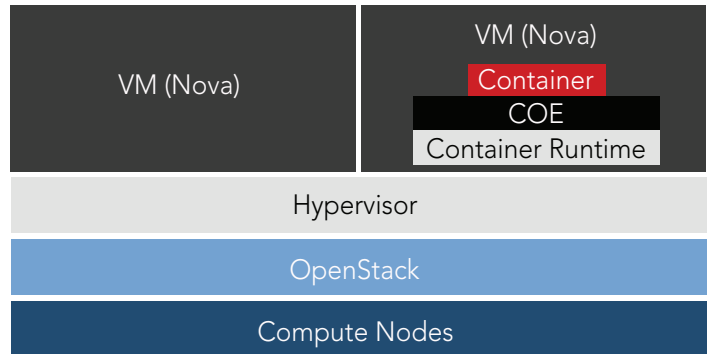
There is substantial investment in hypervisors and VMs today, and these projects cannot simply be scrapped and replaced with a new cloud native approach. CSPs don't have to abandon what they've already built to start deploying cloud native network functions (CNFs). There are many ways to leverage existing investments, but here are some examples of migration paths:

Host containers in VMs: By putting containers in VMs, in a kind of hybrid model, CSPs can take advantage of container features while also leveraging the benefits of their existing architecture of hypervisors and VMs, such as live migration and VM scale and load balancing. The container is essentially treated as a VM, but it can be managed individually and leverages container orchestration engines, such as Kubernetes.


The downsides of this approach are that it creates a complex environment that can be too heavyweight for an edge cloud deployment, while certain performance aspects may suffer.

Bare metal containers in a VM architecture: In contrast to the previous hybrid model, this architecture is more like a dual model. That is, the architecture provides a bare metal environment within OpenStack; the containers and VMs are separated and managed equally. Containers run natively on the bare metal, without a hypervisor, and are managed by container orchestration engines, which improves performance. This is a flexible and proven method for deploying containers and VMs that allows CSPs to maintain their VM legacy investments. But it still has the heavyweight control plane of the OpenStack environment, which is not good for edge cloud deployments.


Containers first with VM support when needed: This architecture essentially flips the previous models around by containerizing OpenStack services and running them in a container on top of a bare metal Kubernetes cluster. The VMs and containers are all managed equally and in the same way through Kubernetes. OpenStack can be run but only when needed to support legacy VMs. This solution is far more lightweight with a thinner control plane, offers better performance, and enables CSPs to deploy the OpenStack containers only when they have VMs that require OpenStack. This architecture is implemented by the StarlingX open source project.



1. Containers in virtual machines (VMs)



2. Bare metal containers in VM architecture



3. "Container first" with VM support when needed

Figure 4. Top options for transitioning to containers at the telco edge

HOW TO BUILD YOUR OWN CLOUD NATIVE EDGE CLOUD WITH STARLINGX

There are many open source projects dedicated to easing the deployment of edge clouds, such as the CNCF, the Linux Foundation's Akraino and EdgeX Foundry projects, and the OpenStack Foundation's AirShip and StarlingX projects.



StarlingX is one of the newest projects, and it stands out because it provides a pre-built distribution for edge clouds that removes much of the complexity of building a cloud native platform. It's a container-based architecture that supports the use of VMs when needed. It addresses some of the diversity and scalability issues of distributed edge cloud deployments with flexibility that is designed for edge and OT use cases. It can scale from a small-footprint single server to a complete mini data center. And it can be architected in a distributed or a centralized model.

StarlingX addresses the key requirements for distributed edge clouds and overcomes the current limitations on OT and edge use cases—namely, zero touch provisioning, workload orchestration, centralized management, edge cloud autonomy, and massive scalability.

Wind River® is focused on accelerating the massive innovation and disruption at the network edge through this important industry initiative. As a commercial deployment of StarlingX, the Wind River Titanium Cloud™ platform will be key to enabling new business opportunities and innovative applications across multiple market segments.

CONCLUSION

Cloud native can help future-proof telco networks at a critical time when the pace of technology changes is constantly increasing. As service providers are formulating strategies for delivering edge, OT, and 5G use cases, they should consider deploying a cloud native container strategy to support distributed edge clouds. Deployments at the edge of the network are far more complex than implementations in large data centers, and there are many different ways to arrive at a cloud native solution while also protecting legacy investments. But with the right distributed edge cloud platform, service providers can cost effectively deliver a wealth of compelling, revenue-generating services.

