The Myth of Cloud Native Workloads

By Olivier Smith
Technical Staff, Office of the CTO

These days, it's virtually impossible to discuss anything telecommunications related without invoking the term *cloud native*. This is true of customer meetings, company workshops, and conversations in most open source communities where I participate.

In Linux Foundation Networking (LFN) in particular, the term cloud native has gained rapid momentum and focus over the past year. A primary reason for this is multiple projects are preparing a new Network Function Virtualization Infrastructure (NFVi) based on Kubernetes. This is significant because Kubernetes is the de facto open source platform for deploying cloud native workloads. As a result, telecommunication service providers will have the ability to deploy 5G cloud native network functions (CNFs) across public, private or hybrid cloud.

This is a potential game changer, as workloads become highly scalable, resilient, distributable, and offer service providers newfound operational agility. But is it really that simple? Do service providers automatically inherit the benefits of cloud native because their infrastructure is using Kubernetes?

Do workloads somehow become cloud native by default? My response is unfortunately, no!

Reaping the benefits of cloud native requires proper design choices in both infrastructure and workload.

Much of the focus in LFN is on the infrastructure itself, to the exclusion of the workloads. For some in the community, specifying the design of workloads is considered "out of scope". For others, an early focus has been on ensuring interoperability and access to infrastructure resources. Neither of these options is sufficient.

Workload Design Matters

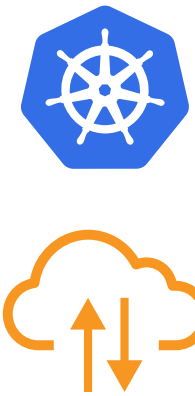

The telecommunication industry's interest in cloud native is driven by the benefits that properly architected workloads achieve. These benefits are realized through use of new design patterns, permitting the workloads to maximize and leverage the capabilities of a distributed cloud infrastructure. Only through such an approach can one, for example, fully leverage Kubernetes to automatically distribute, scale, and heal workloads across a cloud cluster. So what design patterns are imperative? While not an exhaustive list, some of the key architectural design elements should include:

- Use of autonomous business functions addressed as loosely coupled microservices
- API first design for interactions between microservices
- Clear separation and management of stateless and stateful services
- Microservices packaged using lightweight containers
- Deployed using Continuous Integration / Continuous Deployment pipelines
- Container lifecycle orchestrated (i.e. using Kubernetes) to manage and schedule based on demand

The truth is that improperly designed workloads will fail to deliver most, if not all, of the benefits desired of cloud native — even if deployed on a Kubernetes — based infrastructure. For example, a monolithic workload, packaged in a container, might be fully capable of consuming cloud native NFVi resources, but will fail miserably to produce any of the benefits that service providers expect like elasticity and resilience. The point is, being compatible with an infrastructure is not the same as maximizing its capabilities. This is a situation best avoided, but how?

CLOUD NATIVE COMPUTING FOUNDATION

THE LINUX FOUNDATION

3GPP A GLOBAL INITIATIVE

Best Practices for Telco Workloads

As a starting point we should leverage existing best practices and knowledge from the authority on cloud native, the Cloud Native Computing Foundation (CNCF). But a "copy & paste" of these principles and best practices into the telco domain will likely not be sufficient. Instead we must recognize that the telecommunication industry differs from others and work to establish a relevant set of Telecom workload best practices.

So, what are some of the ways in which Telco differs?

Firstly, in the context of the LFN, workloads will run on what effectively is a telco distribution of Kubernetes which supports both CNFs and Virtualized Network Functions (VNFs) that must be capable of discovering and exchanging their capabilities. This is required because service providers have invested heavily in VNFs and will expect new CNFs and existing VNFs to co-exist while transitioning fully to cloud native network functions over time.

Secondly, workloads that will be executing on this new infrastructure are standards-based telecommunication network functions. These must adhere to standards specified by 3GPP to ensure interoperability with other network functions and with the networks of other service providers.

The desire to adopt cloud native must thus be coordinated and aligned with the standards that are critical for the telecom industry.

Recognizing and accommodating for these types of differences is fundamental to maximizing the value of cloud native for service providers. More importantly, such an effort will help advance telecom's cloud native adoption while ensuring existing service provider investments can be leveraged.

Coordination and Collaboration Is Key

One thing is for sure, driving telecom industry adoption of cloud native will require a team effort. Many open source communities have an opportunity to play an important role, including both the CNCF and LFN — and there are signs that we are heading in the right direction.

Within LFN, we are no longer just talking about cloud native, but many of our communities are now embracing this direction in projects and taskforces focused on supporting the next generation of cloud native network functions. We are also seeing changes on a larger scale, such as the recent decision to merge the CNTT and OPNFV into a new joint community named Anuket.

The ambition here is to create a community with tighter coordination and feedback across the design, implementation, and verification processes as LFN progresses its cloud native journey. Similarly, LFN and CNCF are collaborating in different efforts to leverage synergies to support a common goal — wider cloud native adoption.

This brings us to the final aspects of securing cloud native benefits — independent verification. The output of all our efforts including the work of service providers and vendors can be tested and compliance verified via the LFN's new Cloud Native Verification Program (CN OVP). This community is developing a compliance and badging program for both infrastructure and workloads. Importantly, this community plans to verify and badge across several areas from compliance with the new NFVi, CNF orchestration using ONAP, to cloud native compliance. This effort highlights the need for strong collaboration among communities across the Linux Foundation.

If we do this right, our efforts will help service providers gain confidence that verified workloads are compliant with the new infrastructure and more importantly, able to maximize its capabilities.

Better Together

The materialization of performant cloud native workloads is necessary for commercial adoption of cloud native, and equally important to the wider acceptance across all industries. To that end, both LFN and CNCF have strong incentives to continue working together, and even widening collaboration to secure the benefits that properly designed workloads deliver. At the time of writing this piece, a new workgroup within CNCF has been announced — the focus of which will be to set clear guidance and best practices for cloud native telco workloads. This is a fantastic step forward, and I am hopeful that both LFN and CNCF will unite efforts to support an independent verification program that guides service providers and vendors. Together we can take the mystery out of telecom workloads.