Whitepaper
Telecommunications Operators

Cloud iNfrastructure Telco Taskforce

The Cloud iNfrastructure Telco Taskforce (CNTT) reference framework delivers an interoperable cloud infrastructure for telecom services.

Executive Summary

Cloud based or Cloud native applications have long been standard for the Enterprise, but support for Network Functions Virtualization has only recently become a focus in the telecommunications world, as Telecommunications operators seek new ways to reduce costs and allow for network service agility in an increasingly competitive environment. This technological trend started with network functions virtualization (NFV) primarily implemented using OpenStack Cloud infrastructure. Recently, network function containerization leveraging Kubernetes has been gaining attention and interest in telecom networks because it facilitates the adoption of cloud native network functions (CNFs) and even more fine grained services for better scalability and elasticity.

While both workload virtualization approaches are supported by open standards bodies and communities, one aspect that has been slow to reach consensus is the underlying virtualization infrastructure. Commonly called NFVI, these components are an essential element of a virtualized service defining the hardware features and software that are the foundation of virtualization. The lack of a shared common reference model and standards that can be widely used in the industry has slowed adoption across the telecom industry, impacting both the operators and vendors in the ecosystem.

The mission and objective of the Cloud iNfrastructure Telco Taskforce (CNTT) is to develop a framework for standardizing Cloud Infrastructure to increase interoperability between the virtualized workloads and the underlying infrastructure, and to allow for validation and conformance testing. After only a little more than a year, the taskforce is well on its way towards meeting its goals. It has developed a reference model framework for supporting several reference architectures, and has started the process of actualizing these into reference implementations and compliance badging of usable cloud infrastructure.

While there has been amazing progress in such a short time, there is still much work to be done. The rapidly growing CNTT community is encouraging any and all people and companies in the Telecom industry that have an interest in defining a more effective way to build infrastructure to support Cloud Native network functions to join the effort however they can. Some of the areas that need attention include:

- Companies with an interest and resources to do field tests of the reference architectures to validate the same.
- Test engineers to add to the test suites, including helping with the definition of the right tests
- Cloud native expertise and community experience to help guide the direction of the Reference Architectures, Implementations and Conformance workstreams.

Introduction

The CNTT was launched in April 2019 by 10 Telecommunications service providers with the goal of creating a unified reference model for building infrastructure for use in network functions virtualization (NFV) implementations. NFV infrastructure (NFVI) is the hardware and software required to deploy the VNF Components (VNFC-s) of the virtual network functions (VNF) software needed for an NFV-based service. To limit the scope of the project, the decision was made to not include management and orchestration (MANO) which were already well developed Open Source projects in their own right.

CNTT was established to address a perceived gap in the Open Source reference models for building virtualized infrastructure. The situation was that the Telecom operators often had to make hard choices from a variety of proprietary and often incompatible NFVI implementations. This in turn caused friction for vendors and operators in the communities (see Figure 1) because the VNF developer community often had to support multiple infrastructures which led to painful conflicts on which infrastructure to build the code on. On the operator side, the inability to converge on a few referenceable infrastructure platforms has sometimes led to overbuilding and operational inefficiencies. While the challenge started with VM based NFVs, it has been compounded as operators increasingly pivot to other cloud native network functions (CNFs) such as containerized and microservices based architectures as well.

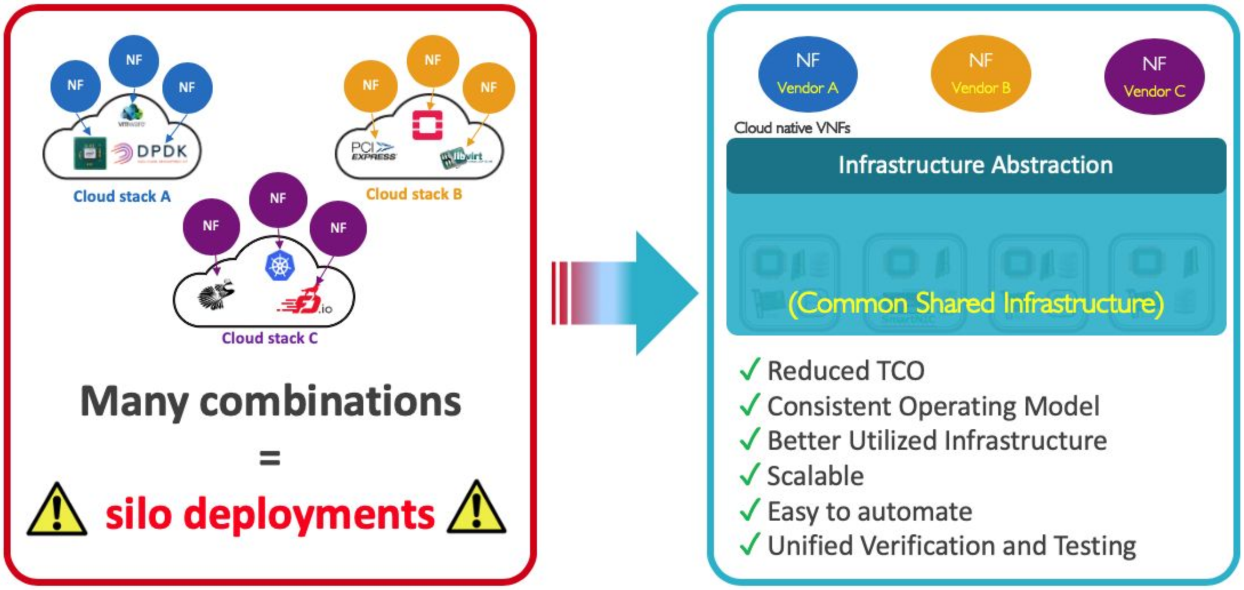

Figure 1. CNTT Problem Statement

Shortly after its initial formation (during the Open Networking Summit (ONS) in San Jose), the CNTT founding members secured the sponsorship and support of both GSMA and The Linux Foundation in Summer 2019, having solicited their sponsorship from the early beginnings. Once the sponsorship was in place, the original operators reached out to the larger industry community including vendors and other related Open Source and Standards projects, and the group membership quickly expanded to telecommunications equipment manufacturers and other vendors to ensure diverse input to the emerging framework. By September 2019, the CNTT workgroup had published its first Reference Model and Reference Architecture for a unified infrastructure framework and presented the work at the Open Networking Summit (ONS) in Antwerp.

CNTT now works in close collaboration with LF Networking (LFN), which facilitates collaboration and operational excellence across open source networking projects, and its testing and integration project, OPNFV. Compliance and verification for infrastructure and cloud native network functions is actualized via the OPNFV Verification Program (OVP). From its inception, CNTT thought testing was needed to augment and validate the development of a new infrastructure reference model specification. Other organizations and projects that work closely with CNTT include European Telecommunications Standards Institute (ETSI) NFV Industry Specification Group, OpenStack Foundation (OSF), Open Network Automation Platform (ONAP), Cloud Native Computing Foundation (CNCF), MEF and TM Forum.

As of the start of 2020, there were approximately 30 CNTT member organizations representing a global group of Operators, telecommunications equipment manufacturers, testing companies, software vendors, and system integrators.

CNTT Framework Overview

CNTT has embraced a scalable, multi-level framework for the infrastructure model and the corresponding set of architectures built on that model. As shown in Fig. 2, the framework is composed of a Reference Model (RM) with the requirement and guidelines that set the direction. Reference Architectures (RAs) provide the specification and guidelines for a particular cloud technology. To date, two RAs have been initiated – one (RA1) is based on OpenStack and the second (RA2) is based on containers using Kubernetes. Reference Implementations (RIs) are developed for each RA with practical information on specifications for installation, lab and resource requirements. The Reference Conformance (RC) provides a mechanism for testing for compliance to the Reference Implementation to ensure interoperability. The Reference Model and the entire framework is open and available to everyone to use for testing and validation.

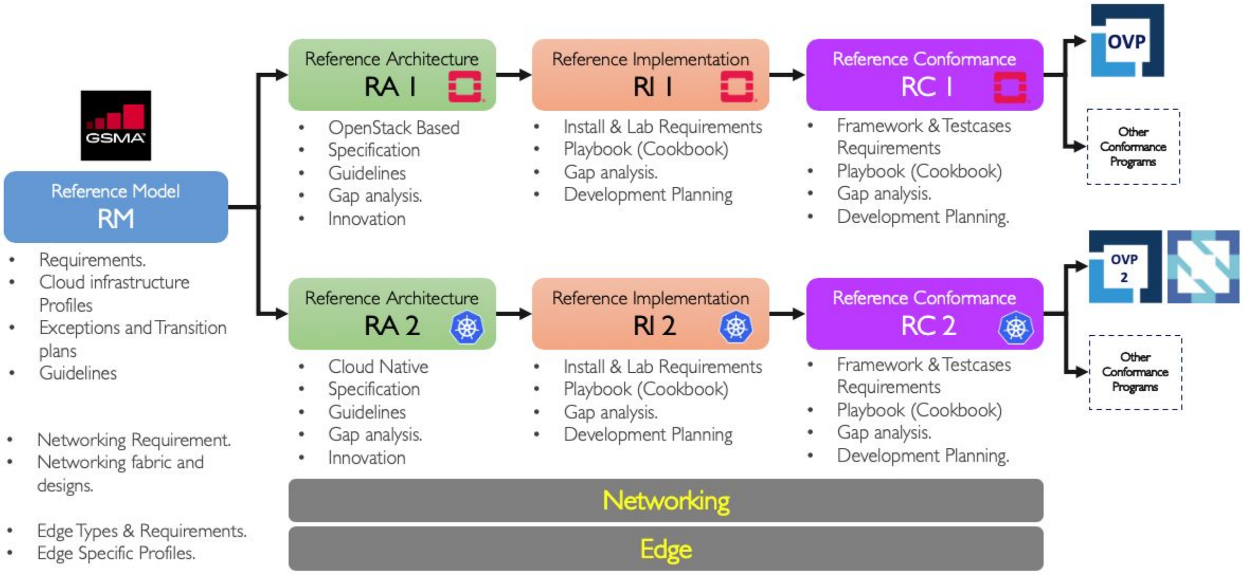

Figure 2. CNTT Reference framework, Workstreams and Focus Groups

One principle for the framework development is to define the Reference Model in such a way that it requires the smallest number of architectures as is practical. To accomplish this, the taskforce has adopted two principles:

- Minimize architecture proliferation by stipulating compatible features be contained within a single Architecture as much as possible.
- Create an additional architecture only when incompatible elements are unavoidable.

To encourage adoption by Operators and vendors alike, these architectures are mandated to first use established technology and systems already common to the industry. Follow-on architectures can utilize new technologies as agreed to by the community members. The task force has built rules to gain commitment from a broad spectrum of the membership before undertaking a new workstream.

Reference Model

The Reference Model (RM) specifies the infrastructure abstraction to provide a uniform, interoperable way for the infrastructure to be compatible with other components outside the scope of the infrastructure, the “actual workloads” if you will, such as the VNFs and CNFs. The intention of the Reference Model is to be as virtualization technology agnostic as possible. The Reference Architecture in turn becomes increasingly specific to a given cloud platform, by profiling hardware and software functionality creating a catalog of capability resources, and interfaces that can be exposed to workloads. The RM specification is a catalog of architecture and requirements that cover:

- Compute
- Storage
- Network
- Acceleration
- Application programming interface (API)
- Multi-tenancy
- Operations and lifecycle management (LCM)
- Security
- Platform and access
- Confidentiality and integrity
- Workload security
- Image security
- Security life cycle management
- Monitoring and security audit
- Standards compliance

An important aspect of the Reference Model is that it specifies the conformance and verification programs and frameworks necessary to ensure a proper implementation of the architecture that can interoperate with the other components of the overall systems.

Another aspect of how the Reference Model was constructed was its ability to allow incorporation of new technologies as they emerge. The CNTT group is well aware that cloud technologies are rapidly evolving, so taking a flexible approach to the model, rather than rigid standards makes more sense for the future robustness and durability of the CNTT foundational principles.

Reference Architecture 1

Reference Architecture 1 (RA1) is based on OpenStack. OpenStack was chosen because it has the advantage of being a mature and widely accepted open source technology; has a strong ecosystem of vendors that support it, the community is managed by the OpenStack Foundation, and, most importantly, it is widely deployed by the global operator community for both internal infrastructure and external facing products and services.

Many of the operators already have existing staff with the right skill sets to support an OpenStack deployment into development, test and production. Another reason OpenStack was chosen for RA-1 is that it has a large active community of vendors and operators, which means that any code or component changes needed to support any newly identified CNTT requirements can be managed through existing project community processes to add and validate the required features through well-established mechanisms.

The May 2020 RA1 release (codenamed Baldy) document relies on the OpenStack Pike release and provides guidance on building a CNTT-conformant cloud infrastructure architecture. It specifies the OpenStack services, features and APIs mandatory that will be tested and verified for conformance as RI-1 and RC-1 are further refined and identified. Beyond the foundational architecture to support reference implementations and conformance testing, RA1 identifies a set of recommended but optional services and features.

Reference Implementation 1

Reference Implementation 1 (RI-1) is a detailed implementation and deployment “cookbook” for CNTT compliant OpenStack deployments. RI-1 includes the detailed requirements for NFVI configurations according to the requirement of the RA and the requirements of each VNF. The RI also includes lab requirements for RI deployment, validation and installer requirements. This reference implementation will include all the capabilities and features defined in the CNTT RM and RA. It is important to note that the primary purpose of the RI is to have a representative implementation of the RA to aid in the development of the associated Reference Conformance test suites.

Details of RI-1 are developed in close coordination with OPNFV to ensure quick deployment and support for the testing frameworks, by utilizing the OPNFV systems, installer automation and testing automation capabilities.

Reference Architecture 2

Reference Architecture 2 grew out of the increasing interest in the support for containerized workloads and services in telecom network infrastructures, and the significant difference between containerized workload infrastructures and VM-based ones. RA-2 describes the high level system components and their interactions, building on the principles and requirements of the CNTT Reference Model and mapping them to real-world containerized components, including Kubernetes (and related) elements. One key aspect of RA-2 is to identify the relevant Kubernetes features and extensions that are best used to support telecom network deployments. Unlike RA1, this work is defining an architecture that is far less mature and widely deployed as telecom network infrastructure. As such, it is expected that there will be more need to interface with the container communities such as CNCF to help with developing new features that are identified as requirements that are needed by the telecom community to support networking workloads. If gaps are identified, the CNTT team can report them back to the Kubernetes community developers for further development.

As of May 2020, RA-2 scope includes:

- Kubernetes capabilities conform to the RM requirements
- Support for CNFs that consist wholly of containers
- Support for CNFs that consist partly of containers and partly of VMs, both of which will be orchestrated by Kubernetes

The Kubernetes RI (RI-2) has also started development of the playbook for testing and validation of the supportability and compatibility with telecom-oriented cloud native network function (CNF) workloads. The timing for the Kubernetes Reference Architecture and Reference Implementation is planned for late 2020.

Reference Conformance

CNTT is using LFN's compliance and verification program to ensure conformance to an RA by providing a framework of test suites and reference data for both VNF developers (and CNF developers in the future).

The Reference Conformance (RC) program builds on work done via the OPNFV Verification Program (OVP), an open source, community-led compliance and verification program that demonstrates the readiness of commercial NFVI and VNFs.

OPNFV developed a framework of testing, and an initial conformance program called OPNFV Verification Program (OVP), where telecom or test lab operators could run the verification tests and earn a badge as OPNFV NFVI or VNF compliant. The OVP established a basic functionality test framework for OpenStack- and VM-based infrastructure. Commercial products adhering to the RM and RA requirements can undergo workload and NFVI compliance testing.

CNTT is developing conformance badging for both software suppliers and testing labs to demonstrate adherence to RA/RM requirements. In addition, the OVP programs are distinct for the cloud infrastructure platform, VNFs, and CNFs, each to their own unique verification requirements. Figure 3 shows the status of the software releases with the Spring 2020 Baldy Release which includes the RA-2 and RC-1 requirements.

CNTT Roadmap

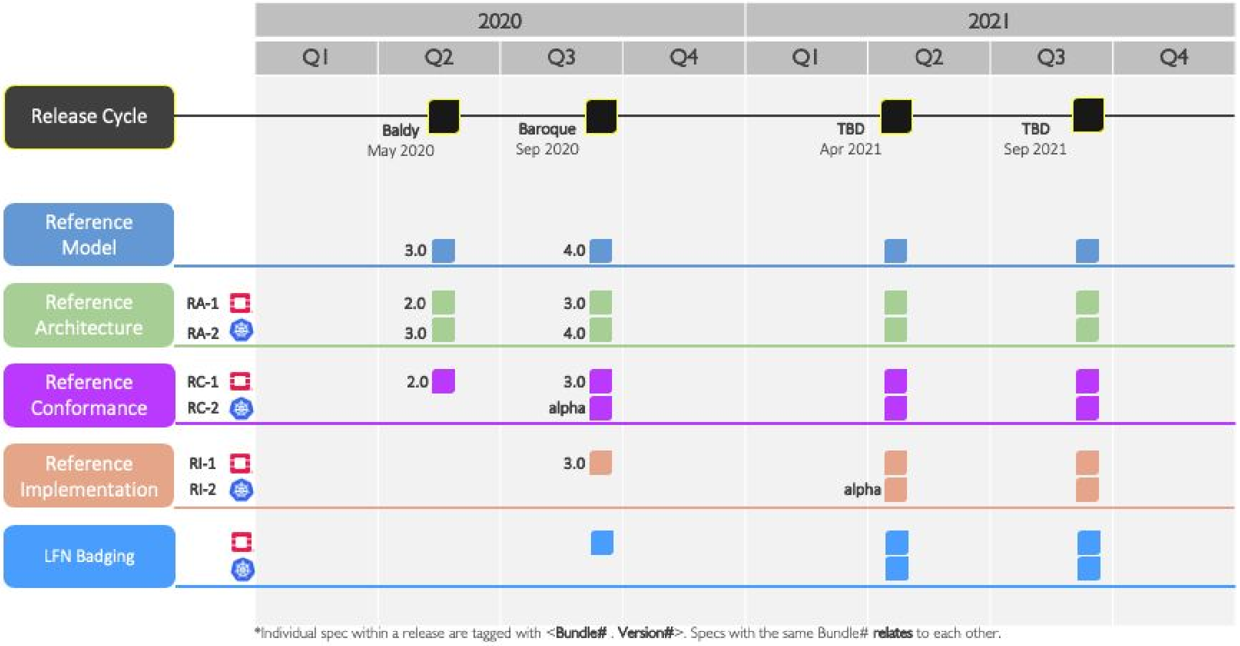

*Individual spec within a release are tagged with <**Bundle#** . **Version#**>. Specs with the same Bundle# **relates** to each other.

Figure 3. CNTT Roadmap

Conclusion

CNTT is developing a common Cloud Infrastructure framework that fills an identified gap in NFV infrastructure development work. Building on existing open source projects and Standards bodies, CNTT is creating an infrastructure framework that will hopefully lead to greater interoperability for both infrastructure and NFV workloads. If successful, it has the potential to reduce the complications and duplicative efforts needed for VNFs software can be assured to work as expected so it can be widely deployed by the telecom community as a whole.

Call to Action

While there has been amazing progress in such a short time, there is still much work to be done. The rapidly growing CNTT community is encouraging any and all people and companies in the Telecom industry that have an interest in defining a more effective way to build infrastructure to support Cloud Native network functions to join the effort however they can. Some of the areas that need attention include:

- Companies with an interest and resources to do field tests of the reference architectures to validate the models
- Test engineers to add to the test suites, including helping with the definition of the right tests
- Cloud native expertise and community experience to help guide the direction of the cloud native workstreams.

For More Information

CNTT GitHub for specifications: https://github.com/cntt-n/CNTT

LF Networking (Linux Foundation): https://www.lfnetworking.org

CNTT Getting Started for New Participants: Community Participant Onboarding - LF Networking - Confluence