

2021-06-09 - Anuket: Multi/Hybrid Cloud

Topic Leader(s)

- Ildiko Vancsa
- Pankaj Goyal
- Beth Cohen
- Walter Kozlowski
- Joao Rodrigues

Topic Overview

60m Ildiko Vancsa (incl. APAC Topic Leader)

Discussing multi- and hybrid cloud requirements and choices as it is driven by edge and other use cases.

The session is a continuation of the similar discussion at the recent PTG: <https://etherpad.opendev.org/p/ecg-ptg-april-2021>

Recording of the PTG discussion: https://zoom.us/rec/share/yMN8QANWzI7FJ_Hh-ys166_sECX2XgSgX-B_jKnnIYh4UjCAv9S32s-9U4mzjUGf.HZIErIBO6LOIMrJz

Relevant RM sections:

- RM 3.5.2 Network Layering and Concepts https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter03.md#3.5.2
- RM 3.6 Storage https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter03.md#3.6
- RM 3.8 Hardware Acceleration Abstraction https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter03.md#3.8
- RM 3.8.3 Workload Placement https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter03.md#3.8.3
- RM Ch 8 Hybrid Multi-Cloud: Data Center to Edge https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter08.md

Relevant RA sections:

- RA1 3.5 Cloud Topology https://github.com/cnnt-n/CNTT/blob/master/doc/ref_arch/openstack/chapters/chapter03.md#3.5
- RA1 4.5 Cloud Topology and Control Plane Scenarios https://github.com/cnnt-n/CNTT/blob/master/doc/ref_arch/openstack/chapters/chapter04.md#4.5

Slides & Recording



YouTube

Please indicate your session type in the blank space below and then remove this Info field.

- Demo / Informational (non-interactive)
 - You may be asked to pre-record this session which will be made available on-demand.
- Live Interactive Session

LFN Staff may elect to publish some videos to YouTube. Please indicate here if you do not want your session to be published to YouTube.

Live Interactive Session

Recording: [Multi-Hybrid Cloud.mp4](#)

Agenda

- Setting the scope for the discussion
 - IaaS? PaaS?
 - Workloads and workload characteristics?
- Definitions - Are these the right definitions?
 - Cloud
 - Multi-cloud environment: clouds from multiple different public cloud providers
 - Hybrid-cloud environment: a mix of private and public clouds
 - Hybrid-multi-cloud: private and multiple public cloud providers
 - Edge - How deeply do we want to go into the edge space? What do we want to/ do not want to define?
 - Edge: EUD, IoT devices
 - From limited to significant compute capacity, storage and networking bandwidth
 - Edge Site: Cloud network periphery or GW
 - Every place where Edge devices connect to the cloud network
 - May consist of limited to significant compute capacity, storage and networking bandwidth nodes for aggregating, processing and filtering data at the edge, including data from dumb devices

- Zero Ops
- Core Site: in a DC, other large sites (IDC, Co-Lo, etc.)
- Edge Cloud: Hybrid multi cloud
 - From single node to small cluster of nodes (HA, resiliency, self-healing, ...)
 - Technology: micro k8s, microstack (single package of all key OSTK services (Keystone, Glance, Nova, Neutron (OVN), Cinder, Horizon, Clustering, monitoring)
- Single Plane of Glass: for seamless management of orchestration of workloads in a hybrid multi-cloud
 - In the last vF2F in February 2021, we discussed the Operator Platform and the need for platform agnostic APIs. Today we want to discuss what other characteristics need to be considered.
 - On June 8, 2021, discussion on Hyperscalars is related (see [2021-06-DD - Anuket: Anuket assured hyperscalars or hyperscaler assured Anuket?](#))
- If discussing workloads
 - Are these the workload characteristics that should be specified when requesting resources?
 - Can schedulers take these characteristics into account when identifying the site, node, .. to host the workload?
 - Workload Characteristics:
 - Resource sizing
 - Storage IOPS
 - Network Latency
 - HW/SW Accelerators
 - Special HW or SW
 - Affinity
 - ...
- Multi-cloud considerations
 - Operational Model
 - Deployment Model (it can facilitate on how the workloads are distributed and configured)
 - Controller suites for Network Services distributed across private and public clouds
 - Challenges: Security, QoS, lack of SLA, access to BSS/OSS systems, orchestration across the networks and hyperscalars.
 - Proposals for multi-cloud support (SDO related)

Minutes

The topic is very broad, so what should we focus on?

- Should focus on workloads and access, not the business models (PaaS, IaaS).
- How can the Telecoms leverage the Hyperscaler resources?

Trends are that vendors are providing environments (Cisco for example) that are not compatible with multi-cloud services. How can telecoms balance the need for control of the infrastructure and the lack of transparency into the hyperscaler architectures?

- Need to rely on brokers and APIs (who controls and develops the APIs). All the hyperscalars have extensive incompatible APIs.
- Open API Forum has been wrestling with the willingness for stakeholders to contribute to OpenAPIs.
 - What is the common interest in open APIs?
 - [Walter Kozlowski](#) thinks that the hyperscalars are in fact interested in working with Telecom operators to define a common set of APIs.
 - GSMA has some work in this area.
 - Do the hyperscalars participate in GSMA?
- [Pankaj Goyal](#) notes that the hyperscalars are in a high growth phase of their business (marketshare acquisition mode).
 - How does 5G fit into this.
 - How do we define the interface with the hyperscalars?
 - Several use cases that need to be addressed – 5G, hosted controller suites, multi-cloud access.
 - Some companies have attempted to bridge public and private cloud (VMWare, is an example) with mixed success.
 - Chapter 6 in RM covers interfaces with Hyperscalars.
 - Needs to better define the ownership and responsibilities.
 - ETSI documents cover the cloud interfaces, where IFA 40 is the stage 2 doc for interfaces and requirements for container solutions. Stage 3 will be profiling (maps stage 2 to the interface in de-facto standards), not an API specification. The profiling for Kubernetes is started in SOL018 document. Question about ETSI's work on container architecture. ETSI is acknowledging the need to cover container and Kubernetes infrastructure, need to have multiple standards.
 - How do we define the management interface to the infrastructure in a multi-cloud environment? RM Table 6-1 https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter06.md#6.2

https://github.com/cnnt-n/CNTT/blob/master/doc/ref_model/chapters/chapter08.md

Focus first on RM and see if we need to create an RA3 (maybe best practices) for how to use the Hyperscalars for Telecom workloads.

Action Items

