

White Paper

Introduction to EdgeStack 2.0

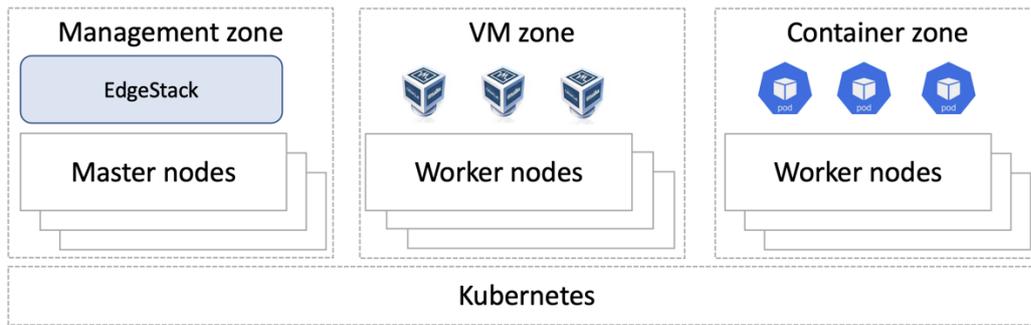
2021. 6. 30

Head of 5GX MEC Product

SK Telecom

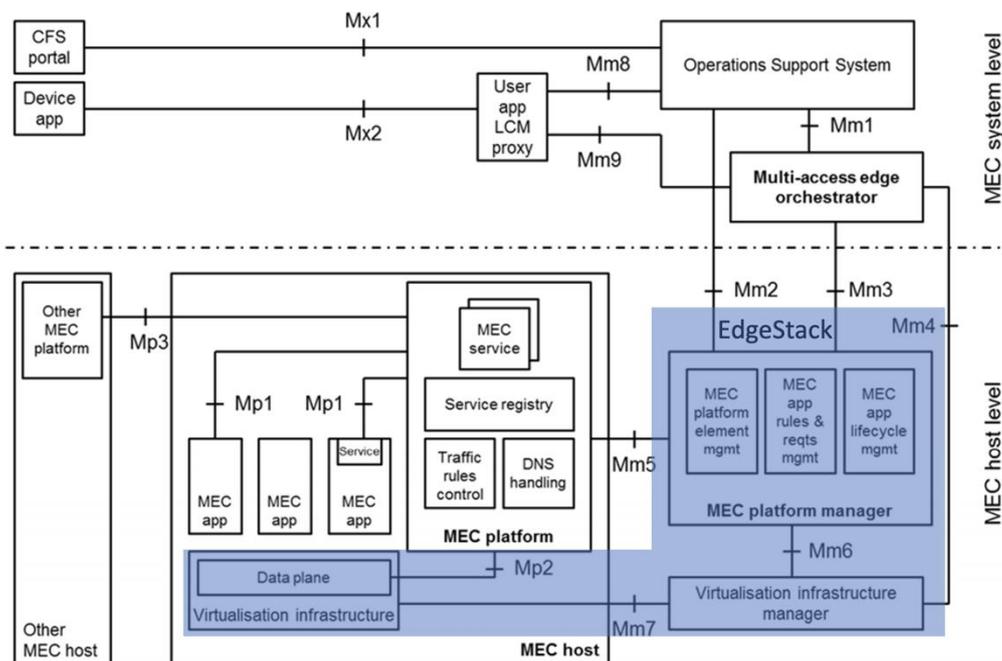
1. Overview

EdgeStack is a cloud stack for **edge environment**, where services are distributed in multiple sites but with smaller scale than central data center. EdgeStack is designed with **lightweight** cloud architecture, which requires **compact** control nodes. However, EdgeStack support **large scale deployment** and **carrier grade high availability** for enterprise level cloud services.



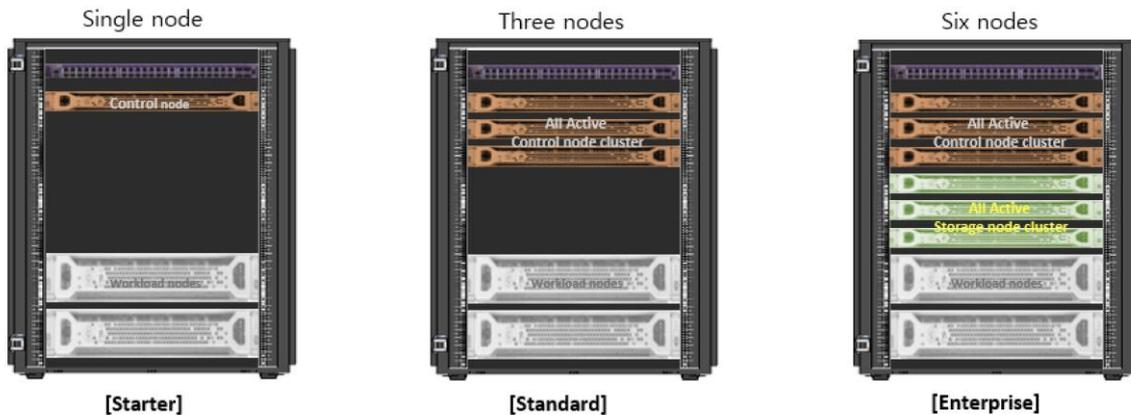
[Fig 1. EdgeStack Deployment Overview]

EdgeStack provides VM and Managed Kubernetes to support both VM and container workloads. They are managed by single Kubernetes technology by utilizing KubeVirt [1] and Cluster-API [2] technologies, which requires less maintenance efforts. Furthermore, we are also planning to integrate the two to single KubeVirt technology in near future. VM workloads and container workloads are deployed and managed in separate hosts called VM zone and container zone, respectively. It is because data plane of each virtual resources is processed in difference ways.



[Fig 2. ETSI Reference MEC Architecture and EdgeStack Coverage]

EdgeStack is compatible with the ETSI MEC Reference architecture [3], and EdgeStack covers the components in blue area of the architecture (refer to Fig 2), which are MEC platform manager, Virtualization infrastructure manager, and data plane in virtualization infrastructure. Following the ETSI MEC architecture, EdgeStack provides REST style MM3 Interface as a north bound interface, which is used to communicate with MEC orchestrator.



[Fig 3. EdgeStack Form Factors]

EdgeStack is deployed on Kubernetes environment and follows Kubernetes HA policy. Therefore, EdgeStack requires three node all active cluster for host level high availability as a standard form factor, but it can be also deployed in a single machine for PoC or testbed, where HA is not mandatory. EdgeStack also supports separate three node all active storage cluster, which provides much reliable storage space and user data is protected more securely. Also, in this case, storage space is expanded as long as the host's physical space is allowed, and we can provide much more storage space to MEC services.

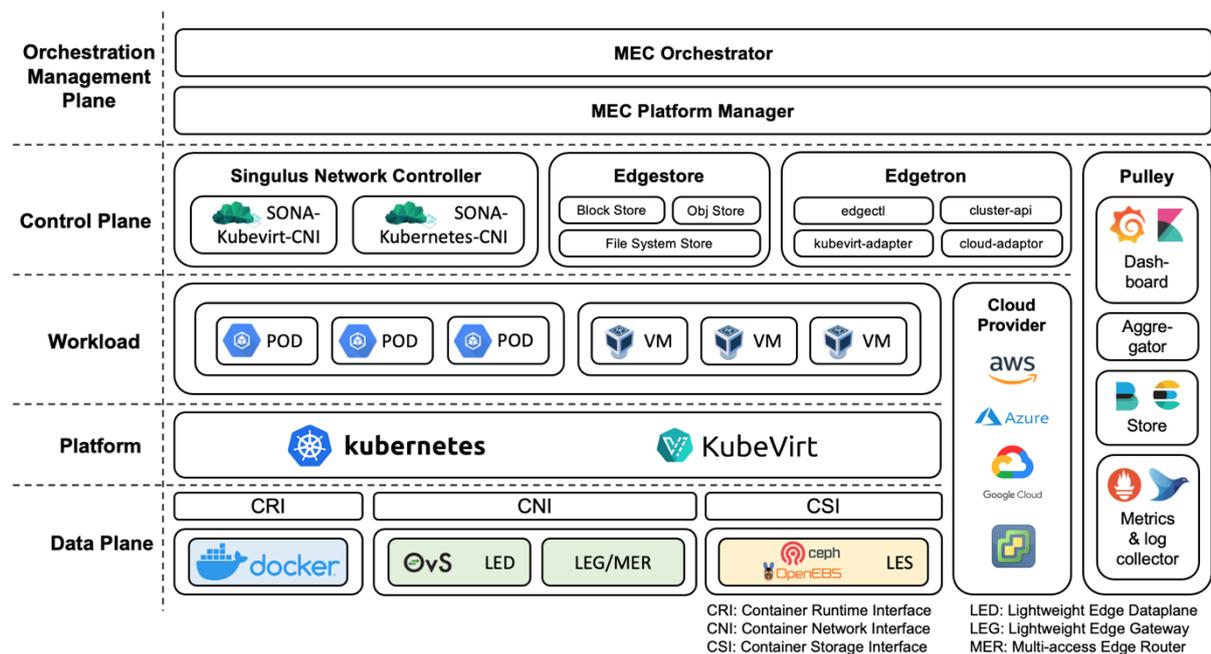
2. Features

EdgeStack provides the following features in version 2.0

- **Large scale VM deployment and life cycle management:** Using KubeVirt technology, VMs are created and managed by pods much more elastically
- **Large scale virtual resource monitoring:** Up to 400 hosts and all of VMs and containers in the hosts are monitored, and the all of the metrics are provided through ElasticSearch DB.
- **High performance virtual overlay network:** Each tenant can create its own overlay network with near line rate throughput using SDN technology
- **Shared block storage:** data are protected with three node cluster based shared storage using Ceph technology [4]
- **Various guest OS support including Windows OS:** Various version of Ubuntu, Centos, Fedora and even Windows are provided as guest OS

- **Integrated system log management:** All systems logs are collected and stored in a central place, and they are stored and managed in ElasticSearch DB [5].
- **MEC Application HA/Load Balance as a Service:** Load balance as a service is provided and can be used for MEC application High Availability
- **Distributed firewall (Security Group):** VMs are protected by security group service. Only traffic defined by security group are allowed to VMs.
- **Scalable lightweight edge gateway:** MEC service is accessed via edge gateway from external network. The edge gateway function is provided as software switch and can be deployed without additional gateway nodes
- **Direct Attached Storage:** Disk volume can be created and managed separately from VMs, and data can be preserved after VM is removed.
- **Public Cloud integration:** AWS EC2 can be created and managed from EdgeStack
- **Managed Kubernetes:** Kubernetes cluster is easily provided on VMs and managed by EdgeStack
- **High performance container network:** Container network performance of the managed Kubernetes is guaranteed as that of Kubernetes on bare-metal and better.
- **VMWare TKG Integration:** VMware TKG cluster can be created and managed from EdgeStack.

3. Architecture



[Fig 4. EdgeStack Architecture]

EdgeStack is using various open sources, and Kubernetes is used as a base system framework. All of the components are containerized and running as pods. In order to provide VMs with more flexibility and in a lightweight way, we use KubeVirt technology, where VMs are provided as pods. In this way, VMs and EdgeStack components are managed in the same manner. To provide managed Kubernetes, Cluster API is used, which is CNCF standard API

to create and manage Kubernetes clusters. Cluster API is used for other third-party cloud solution such as OpenShift and VMware, and they are integrated with EdgeStack without additional development using the Cluster API.

As the network data plane OVS is used, and ONOS SDN controller [6] is used to compute and push flow rules to control the VM and container network flows. As the monitoring system, Prometheus and ElasticSearch are used together with Grafana and Kibana dashboard system.

4. Components

EdgeStack is composed of the following components, and they are developed and managed as separate sub-projects.

EdgeDisk

Edge Disk provides edge storage, a persistent data storage within the MEC platform, providing the storage capacity and enterprise level functions required by each VM/container workload.

EdgeGateway

EdgeGateway is the gateway node that is responsible for north-south traffic routing for VM's and containers. EdgeGateway is basically Open vSwitch-based implementation which is controlled by Kubevirt-CNI . Kubevirt-CNI listens all of the VM, Network, Virtual Router, LB and Floating IP events generated from EdgeStack and programs the gateway flows accordingly based on the EdgeStack network model.

EdgeSpray

EdgeSpray is a deployment automation tool used to deploy the entire EdgeStack software stack to the production environment. Since EdgeStack runs on top of Kubernetes, Kubernetes cluster is required to deploy using kubespray before running EdgeSpray.

Edgetron

Edgetron 2.0 is an infrastructure management component aims to provision not only virtual resource but also containerized resource for edge computing. Edgetron provides northbound interfaces with MEC Platform Manager through ETSI-MM3 standard API, while its southbound interfaces with various cloud providers to ensure the requested resource is properly provisioned. Besides the MM3 API, Edgetron also provides CLI to help the operators and developers to interact with EdgeStack.

Pulley

Pulley is created for monitoring large-scale Kubernetes based virtual resources. Pulley is based on Prometheus, which is widely used in monitoring, and connects with Node-export, k8s-apiserver, kube-state-metrics, fluentd, Kubelet, and kubevirt. For efficient management the collected data are connected to Metricbeat and sent to Elasticsearch for storage. The data stored in Elasticsearch can be used by connecting with Grafana and Kibana and can be used on the

web by extracting them using the Elasticsearch API.

Singulus

Singulus Network Stack, is an SKT in-house built project, implements a set of networking stacks, supports both VM and container workloads using Open Networking Operating System (ONOS) [6] open source SDN controller. Singulus Network Stack aims to provide fast and simple networking stacks, specifically targets to Mobile Edge Computing (MEC) use cases, but not only limited to MEC.

5. Interfaces

As mentioned in Overview, EdgeStack is compatible with ETSI reference architecture and provides MM3 Interface as a north bound interface.

The MM3 interface list is as follows.

[Table 1. EdgeStack MM3 Interface List]

Object	API
VM	/edgetron/resources/kubevirt/vms
VM Action	/edgetron/resources/kubevirt/vms/{name}/action
Image	/edgetron/resources/kubevirt/images
Flavor	/edgetron/resources/kubevirt/flavors
Network	/edgetron/resources/kubevirt/networks
SRIOV Network	/edgetron/resources/kubevirt/sriov_networks
Router	/edgetron/resources/kubevirt/routers
Floating IP	/edgetron/resources/kubevirt/floating_ips
Load Balancer	/edgetron/resources/kubevirt/lbs
Security Group	/edgetron/resources/kubevirt/security_groups
Security Group Rule	/edgetron/resources/kubevirt/security_group_rules
Volume	/edgetron/resources/kubevirt/volumes
Storage Class	/edgetron/resources/kubevirt/storage_classes
Keypair	/edgetron/resources/kubevirt/keypairs
Node (Host)	/edgetron/resources/kubevirt/nodes

Additionally, EdgeStack provides CLI tool, called edgectl. EdgeStack users can enjoy full features of EdgeStack easily with the CLI. The edgectl CLI usage is as follows:

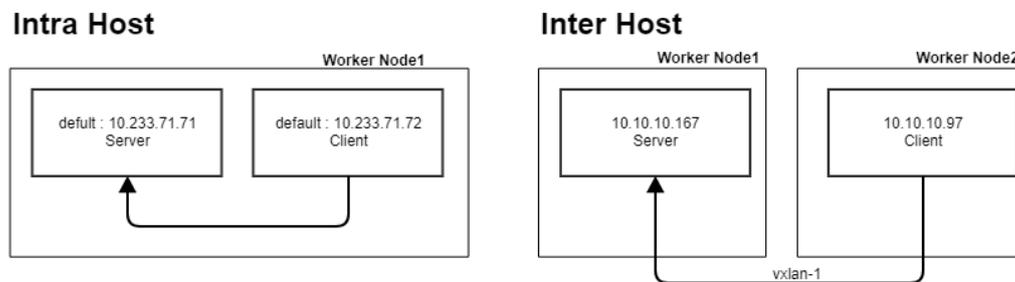
```
[centos@mec-master ~]$ edgectl
associate-fip      list-flavors      remove-image      show-login-users
create-fip         list-images       remove-keypair    show-network
create-flavor     list-keypairs     remove-lb         show-node
create-image      list-lbs          remove-lbr        show-router
create-keypair    list-networks    remove-network    show-sc
create-lb         list-nodes        remove-peer-router show-sg
create-lbr        list-routers      remove-router     show-vm
create-network    list-sces         remove-sg          show-vm-events
create-peer-router list-sgs          remove-sgr         show-vm-logs
create-router     list-sriov-intfs remove-sriov-network show-vm-manifest
create-sg         list-sriov-networks remove-sriov-policy show-vol
create-sgr        list-sriov-pfs   remove-vm         start-vm
create-sriov-network list-sriov-vfs  remove-vol        stop-vm
create-sriov-policy list-vms          restart-vm        test-ping
create-vm         list-vols         show-flavor       unpause-vm
create-vol        network-cli       show-guest-info   update-lb
disassociate-fip  pause-vm          show-image        update-router
edit-vm-sg        remove-fip        show-keypair
list-fips         remove-flavor     show-lb
```

[Fig 5. edgectl command list]

6. Performance

We have measured the network performance of EdgeStack and the results are shown in Table 2 and Table 3.

The Table shows the layer 3 performance test using iperf tool [7]. As we can see in Fig. 6, network through between two VMs created by EdgeStack with two scenarios, inter-host and intra-host.



[Fig 6. Layer 3 network performance test scenarios]

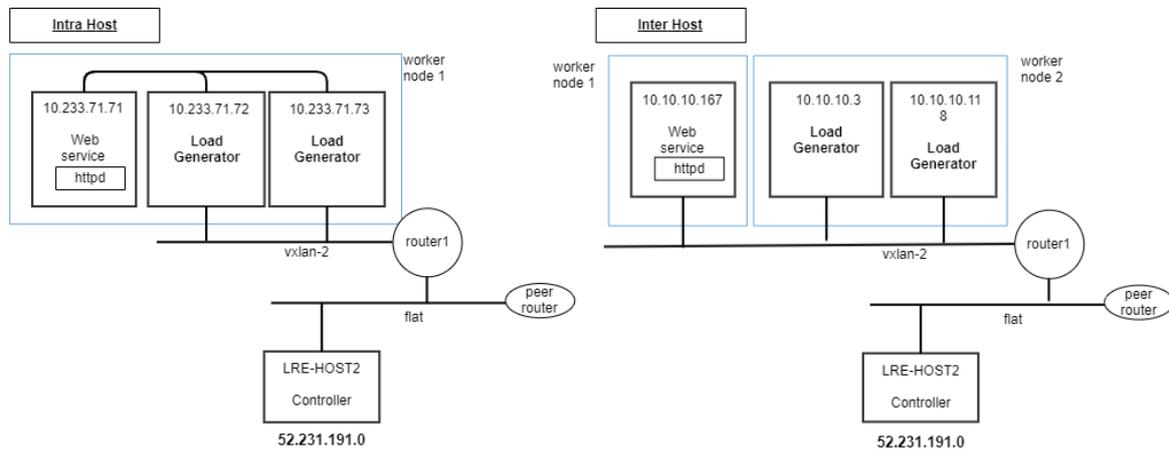
[Table 2. Network performance test with iperf]

MTU Size	128B	256B	500B	1000B	1500B	2000B
Intra Host (Gb/s)	15.3	15.4	14.5	14.9	14.5	14.1
Inter Host (Gb/s)	1.01	2.79	5.80	6.44	7.40	7.40

According to the test, we have confirmed that with 1,500B MTU size, which is the general MTU size, the performance reaches up 7.4 Gbps on average (inter-host case) without any

data plane acceleration technique. We will further improve the performance with various data plane acceleration techniques such as SR-IOV[8] and DPDK[9] in near future.

Additionally, we have performed the application-level network performance test using HTTP traffic, which is dominant traffic in Internet. We tested with upload and download scenarios with various target file sizes.



[Fig 7. Application level network performance test scenarios]

[Table 3. Application level network performance test with HTTP]

	Traffic	File Size (MB)	TPS	Throughput (Gb/s)	Response Time (sec)
Intra Host	Download	10	119.66	10.07	0.19
		500	3.19	13.44	3.09
		1,000	1.60	12.79	5.24
	Down(50%) / Up(50%)	100/10	13.68/13.96	12.74	0.86/0.32
Inter Host	Download	10	83.58	7.51	0.55
		500	2.02	8.74	5.73
		1,000	1.05	8.73	10.8
	Down(50%) / Up(50%)	100/10	9.38/9.97	8.77	1.00/0.27

According to the test, we found out that the total application throughput with HTTP traffic goes up to 8.7 Gbps, which is slightly higher than iperf network test. We presume that it is because the statistical bias from sum of many HTTP clients within spacious response time.

Usages

EdgeStack is outcome of commercial MEC product development project. EdgeStack is used as edge cloud solution of Generic MEC, Media MEC, and Vision-AI MEC. MEC Orchestrator is built on top of EdgeStack using MM3 API, and Biz Portal, which is customer face portal and includes business logic of Vision-AI and Media MEC, is also built on top of MEC Orchestrator communicating with MM1 API.

7. References

- [1] KubeVirt: <https://kubevirt.io/>
- [2] Cluster-API: <https://cluster-api.sigs.k8s.io/>
- [3] ETSI MEC Reference Architecture:
https://www.etsi.org/deliver/etsi_gs/mec/001_099/003/02.01.01_60/gs_mec003v020101p.pdf
- [4] Ceph: <https://docs.ceph.com/en/latest/>
- [5] ElasticSearch: <https://www.elastic.co/kr/elasticsearch/>
- [6] ONOS: <https://opennetworking.org/onos/>
- [7] iperf: <https://github.com/esnet/iperf>
- [8] SR-IOV: https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux_openstack_platform/7/html/networking_guide/sr-io-v-support-for-virtual-networking
- [9] DPDK: <https://software.intel.com/content/www/us/en/develop/articles/open-vswitch-with-dpdk-overview.html>

Authors: 김시리, 남소원, 박민철, 박병수, 변상윤, 신상호, 이건, 장효주