

# ClusterSlice: A Zero-touch Deployment Platform for the Edge Cloud Continuum

Lefteris Mamas (Member, IEEE), Sotiris Skaperas (Member, IEEE), Ilias Sakellariou (Member, IEEE)  
Department of Applied Informatics, University of Macedonia, GR-54636 Thessaloniki, Greece  
{emamas, sotskap, iliass}@uom.edu.gr

**Abstract**—We demonstrate ClusterSlice, an open-source solution for automated Kubernetes-center deployments for the edge continuum. ClusterSlice is an infrastructure-as-a-service, platform-as-a-service, and application-as-a-service solution, supporting: (i) declarative deployment slice definitions; (ii) infrastructure-on-demand capabilities over multiple heterogeneous domains; (iii) composable Kubernetes deployments, supporting multi-clustering as well as various Kubernetes flavors and intra-cluster/inter-cluster network plugins; (iv) configurable application deployment; and (v) experimentation automation.

**Index Terms**—Kubernetes, Edge Computing, Edge Cloud System, Zero-touch Deployment, Cloud-Network Slicing.

## I. INTRODUCTION

ClusterSlice is an open-source Kubernetes-centered Edge Cloud solution, able to convert compute resources from bare-metal or hypervisor-level to fully-operatable cloud-network slices. One of its stronger features is the introduction of well-designed abstractions, that reduce deployment complexity with improved reliability and reproducibility, configuring the corresponding slice parts in a parallel fashion and zero-touch manner. ClusterSlice is cloud-native and fully integrated to Kubernetes, a widely used orchestration facility with impressive automation, fault-tolerance, scalability, resource optimization capabilities and, most importantly for the context of this work, a highly modular architecture. The latter allows Clusterslice to implement its functionality by harvesting on the concepts of Custom Resource Definitions (CRDs) and the Operator pattern, and to be fully managed via the `kubectl` tool. Due to this tight integration, it benefits from all Kubernetes reliability features. ClusterSlice supports:

- *Declarative automated operation*: Slice definition is done at a higher level in the form of YAML files, and includes compute resources to be utilized, Kubernetes configuration, and the application modules to be installed.
- *Infrastructure-as-a-service capabilities*: ClusterSlice supports the utilization of heterogeneous physical and virtual resources through a common abstraction, including open testbed infrastructures (e.g., CloudLab), as well as Virtual Machines (VMs) in XCP-ng and VirtualBox virtualization systems.
- *Platform-as-a-service features*: It supports multiple Kubernetes flavors, such as vanilla Kubernetes (k8s), k0s, k3s, and microk8s, as well as network plugins for both intra-cluster (e.g., Flannel, Calico, Kube-ovn, etc.) and inter-cluster communication (e.g., Submariner).

- *Application-as-a-service attributes*: The slice definition supports deployment of configurable application modules, i.e., the definition of applications to deploy, k8s extensions and modular OS configurations, ensuring maximum compatibility with heterogeneous systems.
- *Multi-cluster and multi-domain capabilities*: ClusterSlice operation can span across multiple heterogeneous deployment environments through technology-specific infrastructure managers, establishing multi-cluster operation and communication, such as Ligo, OCM, or Submariner.
- *Experimentation automation*: ClusterSlice is enabled advanced, automated deployment and management of experiments, based on its characteristics of reproducibility in the deployment of slices, and actively supported automation capabilities, such as comparable evaluations of network plugins and anomaly detection workflows.

ClusterSlice provides infrastructure-, platform- and application-as-a-service capabilities over the edge continuum. Relevant infrastructure control solutions include SLICES (<https://portal.slices-sc.eu>) and FABRIC [1] test-beds, which enable multiple infrastructure federations to inter-operate (e.g., Emulab, CloudLab, and others). Their novel resource control features and APIs (e.g., jFed, python APIs, AMs, etc.) make them re-usable for higher-level deployment abstractions. For example, CloudNativeLab (<https://practicum.cloudnativelab.ilabt.imec.be>) enables the creation of a Kubernetes cluster on SLICES testbeds, using a defined resource configuration for the cluster. Furthermore, 5G-CDN [2] and NECOS [3] are slicing infrastructures capable of allocating resources across multi-domain infrastructures.

The platform- or application-oriented solutions, include EdgeNet [4], SLATE [5] and Kubernetes-native testbed (<https://github.com/kubernetes-native-testbed>). EdgeNet [4] is a platform-as-a-service solution over multiple, crowd-sourced k8s worker nodes, implementing an edge continuum through CRDs and operators. SLATE federates science platforms, allowing sites to delegate service deployment and configuration to designated application administrators. The Kubernetes-native testbed integrates various microservices-based applications, CI/CD environments, and monitoring tools.

ClusterSlice supports multiple infrastructure managers under a common design abstraction, so all above mentioned solutions could be jointly utilized and federated. Furthermore,

it also supports both platform-as-a-service capabilities (e.g., composable Kubernetes clusters or multi-cluster deployments) and application-as-a-service capabilities, defining the application or Kubernetes extensions to use: all at the level of a declarative slice definition.

In the rest of sections, a brief presentation of ClusterSlice is covered in Section II, whereas sections III and IV detail the demo run and conclude the paper, respectively.

## II. PROPOSED SYSTEM

The automated transformation of resources from bare-metal or hypervisor setups into fully-operational Kubernetes-based edge cloud deployments, follows a unique approach: it benefits from k8s capabilities based on novel design abstractions.

ClusterSlice design is empowered by the Kubernetes paradigm, and its implementation heavily utilizes Custom Resource Definitions and Kubernetes Operators, thus inheriting the *reliability*, *scalability*, and *resource optimization* capabilities of Kubernetes. The implementation carefully avoids the use of APIs or technology-specific software, relying only on common tools such as *bash*, command-line utilities, and SSH through Ansible. This preserves *simplicity*, *portability*, and *composability* in line with the Unix philosophy.

Furthermore, ClusterSlice introduces *innovative abstractions* on different levels of its architecture:

- *node and cluster level automation abstractions*, i.e., Resource Managers (RMs) and Slice Operators (SOs), respectively. These act like digital twins of compute resources (i.e., physical hosts or VMs) and k8s clusters. ClusterSlice allocates and declaratively configures such containers, which, in turn, represent and actively configure the resources they represent.
- *infrastructure manager abstractions* that interface with diverse cloud/virtualization systems (e.g., VirtualBox or XCP-ng) and even testbed facilities (e.g., CloudLab), maximizing *compatibility* across heterogeneous systems and seamless *future-proof* integration.
- *multi-cluster and multi-domain abstractions*, to create and oversee multiple clusters and deploy the respective software, e.g., Ligo, Karmada or OCM, in the form of application modules.

A more detailed description of the ClusterSlice architecture and its individual components can be found in [6]. One of ClusterSlice’s strongest features lies in its ability to deploy a diverse range of characteristics for the deployment slice through a declarative definition. This is further demonstrated in the following section.

## III. DEMO DESCRIPTION

The demo involves the preparation of various slice deployment descriptors (e.g., as specified by the demo audience), in the form of YAML files. Subsequently, these descriptors are used for deploying corresponding cloud-network slices. Here, we provide a specific example involving three Kubernetes clusters across heterogeneous domains, accompanied by a multi-cluster management system, specifically Ligo. The

scenario considers the following diverse set of deployment environments:

- A local two-server test-bed hosting the XCP-ng cloud environment, known as “swntestbed”.
- A workstation PC hosting VirtualBox, named “lefteris”.
- The CloudLab test-bed, spread across various locations in the USA, referred to as “cloudlab”.

The declarative definition of the deployment is illustrated in Fig. 1. In the YAML file, we present the overarching properties of the experimentation slice, i.e., the name, namespace, server deployment strategy, and user credentials. After that, the specification of the three clusters follows. Each one of these specifications includes information on node, Kubernetes, and application requirements:

- *Cluster node requirements* such as the number of nodes in each cluster, the OS image to be deployed in master and worker nodes, and the node types. It should be noted that each cluster possesses its own deployment domain and unique resource type, for instance `pc3000` physical server at the CloudLab cluster or VM in the XCP-ng or VirtualBox clusters.
- *Kubernetes cluster configuration* supporting the full range of versions, i.e., vanilla, k0s, k3s, etc., and multiple network fabrics, e.g., Flannel, Calico, KubeRouter, etc., allowing for any (valid) combination of the previous, include requesting a specific version of the above, enhancing the reproducibility aspect of the deployment.
- *Application requirements*, that for this scenario include the Ligo Kubernetes extension. It is interesting to note that the Ligo master deployment (cluster named `liqo`), specifies, in a rather declarative manner, the participating clusters in the multi-cluster deployment (`'peers': '[liqo1,liqo2]'`).

The deployment is initiated by issuing a single `kubectl apply` command. The first phase involves allocation of resources for the three clusters, which is handled by the appropriate infrastructure managers depending on their type; for instance in the case of CloudLab by the *Testbed Infrastructure Manager* whereas in the case of VMs, an embedding process might also occur which leads to VMs allocated to specific cloud servers, via the *Cloud Infrastructure Managers*.

During the second phase, the deployment process involves several steps, including: (a) OS image installation and configuration, (b) Kubernetes cluster setup for the three distinct clusters (`liqo`, `liqo1`, and `liqo2`), and (c) the establishment of the Ligo multi-cluster management system, in this specific scenario.

After completing the above phase, Ligo is successfully implemented across three diverse clusters, as illustrated in Fig.2. At any deployment stage, executing the same command provides the status of all involved compute resources. Multiple such examples will be available during the demo.

## IV. CONCLUSIONS

ClusterSlice offers a declarative, effortless, zero-touch solution for transforming testbed resources, even from bare-metal,

```

apiVersion: "swn.uom.gr/v1"
kind: MultiClusterSliceRequest
metadata:
  name: liqo
spec:
  name: liqo
  namespace: swn
  deploymentstrategy: balanced
  credentials:
    username: clusterslice
    password: sha-512-encoded-password
  clusters:
    - name: liqo
      deploymentdomain: swntestbed
      infrastructure:
        masters:
          count: 1
          osimage: "ubuntu-22-clean"
          mastertype: "vm"
        workers:
          count: 1
          osimage: "ubuntu-22-clean"
          workertype: "vm"
      kubernetes:
        kubernetestype: "vanilla"
        networkfabric: "flannel"
      applications:
        - name: liqo-master
          scope: cluster
          parameters: "{peers:[liqo1,liqo2]}"
    - name: liqo1
      deploymentdomain: lefteris
      infrastructure:
        masters:
          count: 1
          osimage: "ubuntu-22-clean"
          mastertype: "vm"
        workers:
          count: 1
          osimage: "ubuntu-22-clean"
          workertype: "vm"
      kubernetes:
        kubernetestype: "vanilla"
        networkfabric: "flannel"
      applications:
        - name: liqo-peer
          scope: cluster
          sharefile: "liqo1-peer-join.sh"
    - name: liqo2
      deploymentdomain: cloudlab
      infrastructure:
        masters:
          count: 1
          osimage: "UBUNTU22-64-STD"
          osaccount: "lmamatas"
          mastertype: "pc3000"
        workers:
          count: 1
          osimage: "UBUNTU22-64-STD"
          osaccount: "lmamatas"
          workertype: "pc3000"
      kubernetes:
        kubernetestype: "vanilla"
        networkfabric: "flannel"
      applications:
        - name: liqo-peer
          scope: cluster
          sharefile: "liqo2-peer-join.sh"

```

Fig. 1: A Multi-Cluster & Multi-Domain Slice Descriptor

```

user@boss:~/clusterslice/examples$ kubectl get computeresources

```

NAME	IP	RESOURCETYPE	STATUS	DOMAIN	SLICE	APPS
boss	195.251.209.196	bossmaster	up	swntestbed		
cloudlabutah		testbed	up	cloudlab		
coralone	195.251.209.192	cloudserver	up	swntestbed		
coraltwo	195.251.209.193	cloudserver	up	swntestbed		
fs	195.251.209.198	bossworker	up	swntestbed		
kubem1	195.251.209.228	mastervm	allocated	swntestbed	liqo	liqo-master
kubem2	195.251.209.229	mastervm	free	swntestbed		
kubem3	195.251.209.231	mastervm	free	swntestbed		
kubew1	195.251.209.223	workervm	allocated	swntestbed	liqo	
kubew2	195.251.209.224	workervm	free	swntestbed		
kubew3	195.251.209.225	workervm	free	swntestbed		
kubew4	195.251.209.226	workervm	free	swntestbed		
kubew5	195.251.209.227	workervm	free	swntestbed		
lefteris	195.251.209.153	cloudserver	up	lefteris		
npc1	195.251.209.203	mastervm	allocated	lefteris	liqo1	liqo-peer
npc2	195.251.209.204	workervm	allocated	lefteris	liqo1	
npc3	195.251.209.205	workervm	free	lefteris		
ops	195.251.209.197	edgenet	up	swntestbed		
pc308-pg-utah	155.98.39.108	masternode	allocated	cloudlab	liqo2	liqo-peer
pc340-pg-utah	155.98.39.140	workernode	allocated	cloudlab	liqo2	

Fig. 2: Compute Resources after the Slice Deployment

into fully operational Kubernetes slices. This is achieved via well-designed abstractions and heavy exploitation of the Kubernetes CRD and Operator patterns.

The software is available as open source via its GitHub repository (<https://github.com/SWNRG/clusterslice>), that also includes a video of the demo presented above. We have run multiple experimentation scenarios, as presented in [6], deploying a total of 410 clusters and 1530 nodes, without a single failure, an initial strong indication of the reliability advantages of ClusterSlice.

## ACKNOWLEDGMENTS

This research is part of the project “Experimental Assessment Of Network And Compute Resource Optimization Methods (NCOPT)”, that has received funding from University of Macedonia Research Fund.

## REFERENCES

- [1] I. Baldin, A. Nikolich, J. Griffioen, I. I. S. Monga *et al.*, “Fabric: A national-scale programmable experimental network infrastructure,” *IEEE Internet Computing*, vol. 23, no. 6, pp. 38–47, 2019.
- [2] P. Valsamas, I. Sakellariou, S. Petridou, and L. Mamatas, “A multi-domain experimentation environment for 5G media verticals,” in *Proc. IEEE Int. Conf. Comput. Commun. Workshops (INFOCOM WKSHPs)*, Paris, France, Apr. 2019, pp. 461–466.
- [3] S. Clayman, A. Neto, F. Verdi, S. Correa *et al.*, “The NECOS approach to end-to-end cloud-network slicing as a service,” *IEEE Commun. Mag.*, vol. 59, no. 3, pp. 91–97, 2021.
- [4] B. C. Şenel, M. Mouchet, J. Cappos, O. Fourmaux *et al.*, “Edgenet: a multi-tenant and multi-provider edge cloud,” in *Proc. 4th ACM Int. Workshop Edge Syst., Analytics and Netw. (EdgeSys)*, Edinburgh, Scotland, UK, 26 Apr. 2021, pp. 49–54.
- [5] J. Breen, L. Bryant, G. Carcassi, J. Chen *et al.*, “Building the slate platform,” in *Proc. ACM Pract. Experience Adv. Res. Comput. (PEARC ’18)*, Pittsburgh, PA, USA, Jul. 2018, pp. 1–7.
- [6] L. Mamatas, S. Skaperas, and I. Sakellariou, “Clusterslice: Slicing resources for zero-touch kubernetes-based experimentation,” University of Macedonia, Thessaloniki, Greece, Tech. Rep., 2023. [Online]. Available: <https://ruomo.lib.uom.gr/handle/7000/1757>