

Cross-Cluster Networking to Support Extended Reality Services

Theodoros Theodoropoulos, Luis Rosa, Abderrahmane Boudi, Tarik Zakaria Benmerar, Antonios Makris, Tarik Taleb, *Senior Member, IEEE*, Luis Cordeiro, Konstantinos Tserpes and JaeSeung Song *Senior Member, IEEE*

Abstract—Extended Reality (XR) refers to a class of contemporary services that are intertwined with a plethora of rather demanding Quality of Service (QoS) and functional requirements. Despite Kubernetes being the de facto standard in terms of deploying and managing contemporary containerized microservices, it lacks adequate support for cross-cluster networking, hindering service-to-service communication across diverse cloud domains. Although there are tools that may be leveraged alongside Kubernetes in order to establish multi-cluster deployments, each one of them comes with its own drawbacks and limitations. The purpose of this article is to explore the various potential technologies that may facilitate multi-cluster deployments and to propose how they may be leveraged to provide a cross-cluster connectivity solution that caters to the intricacies of XR services. The proposed solution is based on the use of two open-source frameworks, namely Cluster API for multi-cluster management, and Liqo for multi-cluster interconnectivity. The efficiency of this approach is evaluated in the context of two experiments. This work is the first attempt at proposing a solution for supporting multi-cluster deployments in a manner that is aligned with the requirements of XR services.

Index Terms—Kubernetes, Cloud, Edge, Continuum, Cluster API, Liquid Computing, Liqo, Network, XR, Immersive Services, 5G, and 6G.



1 INTRODUCTION

Contemporary applications are deployed in the form of containerized microservices. To facilitate the underlying orchestration complexity of containerized microservice deployments, the notion of container orchestration frameworks was introduced. Kubernetes [1] is an extensible open-source orchestration platform for automating software deployment, scaling, and management of containerised workloads and services and is considered to be the standardized way of orchestrating containers and deploying distributed applications. While Kubernetes is extremely popular in cloud computing environments, lightweight versions, such as K3S, are often deployed in Edge computing environments.

Furthermore, contemporary applications, such as eXtended Reality (XR), are often intertwined with a plethora of

demanding Quality of Service (QoS) and functional requirements. The backbone of XR applications relies on providing an immersive experience for the various end-users. Providing acceptable levels of immersion requires extremely low latencies and high bandwidths. The scientific literature has showcased that for an end-user experience to be considered satisfactory, the end-to-end latency shall not be greater than 15ms, and the available bandwidth should be scalable up to 30 Gbps [2]. Furthermore, XR applications are extremely demanding in terms of computational resources since they incorporate functionalities such as the rendering of complex 3D models and the use of highly-defined graphics.

Another important requirement of XR services is the need for end-user equipment to be as light-weight and inexpensive as possible. While cloud computing can shift the computational adequacy burden to various remote resources, thus allowing end-user devices to be mobile and cost-effective, it cannot fully support immersive applications that require low latency and high bandwidth since the end-user devices are usually far from the cloud servers. This fact leads to processing and network overheads, thus resulting in low performance and high latency. Edge computing aims at reducing the amount of data that needs to be transmitted to remote clouds and allows for data processing near the data sources. Thus, edge computing can provide faster response times, higher transfer rates, and better scalability and availability. Consequently, running XR services in a distributed fashion across the cloud-edge fabric would benefit application developers and help keep up with the aforementioned QoS requirements [3].

Such an XR service deployment scenario that spans across the cloud-edge fabric is depicted in Fig. 1, whereby some distant users must collaborate in a virtual reality en-

-
- *Theodoros Theodoropoulos, Antonios Makris, and Konstantinos Tserpes are with the Department of Informatics and Telematics, Harokopio University of Athens, 17778 Tavros, Greece and with the School of Electrical and Computer Engineering, National Technical University of Athens, 15773 Zografou, Greece*
 - *Luis Rosa and Luis Cordeiro are with OneSource, 3030-384 Coimbra, Portugal*
 - *Abderrahmane Boudi, and Tarik Zakaria Benmerar are with ICTFICIAL OY, 02130 Espoo, Finland*
 - *Tarik Taleb is with the Electrical Engineering and Information Technology, Ruhr University Bochum(RUB), 44801 Bochum, Germany, and also with the Department of Convergence Engineering for Intelligent Drone, Sejong University, Seoul 05006, South Korea*
 - *JaeSeung Song is with the Department of Convergence Engineering for Intelligent Drone, Sejong University, Seoul 05006, South Korea.*

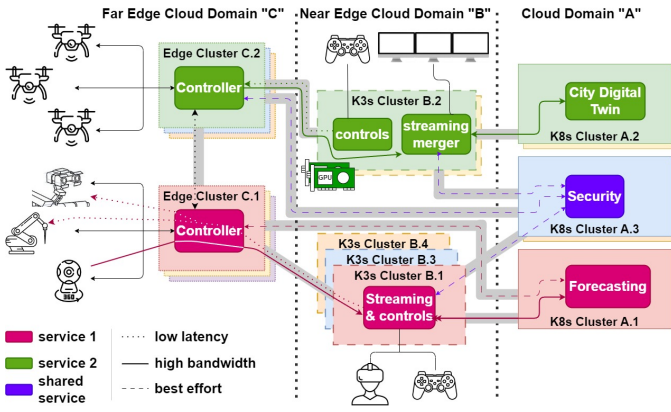


Fig. 1. Multi-Cluster Example Use Case.

environment and would require many services to run concurrently over a wide area. As stated before, XR services require extremely low latency, which means that part of the service must be deployed close to stakeholders (i.e., domains B and C in Fig. 1). In addition, by having many users collaborate closely, communications paths need to be spun between the deployed services in a peer-to-peer fashion; whilst also communicating with some carefully placed shared services, e.g., used for security or synchronization purposes. Other resource-intensive services can also be placed in the cloud. It is clear that for this relatively simple, yet realistic, XR service, and with the many shifting requirements of the sub-services and their placement, leveraging cloud-edge deployments would greatly support the provisioning of such services.

Aside from XR services, cross-cluster management and connectivity are pivotal in diverse domains such as wireless networks, edge computing, 5G applications, Service Function Chains (SFCs), and the scheduling of containerized workflows. Effective cross-cluster strategies enable seamless communication and coordination across disparate systems, ensuring optimal resource utilization and enhanced performance. In wireless networks and 5G, cross-cluster connectivity supports robust, low-latency communication essential for real-time applications. In edge computing, it enables efficient data processing and service delivery by integrating edge nodes. Similarly, for containerized workflows, cross-cluster orchestration ensures scalability, and efficient workload distribution across multiple clusters, leading to improved operational efficiency. Finally, in the frame of SFCs, it enables the chaining of network services across different clusters, facilitating complex service delivery across distributed environments.

Although it is technically feasible for a single Kubernetes cluster to span across various edge & cloud sites that may be at vastly different locations, multi-cluster deployments offer numerous advantages. These include the ability to shift workloads between clusters to prevent congestion or failure, minimizing downtimes by seamlessly transferring workloads to alternative clusters, overcoming size constraints of single-cluster deployments, ensuring compliance with country/region-specific regulations for data storage in cloud-based workloads, allowing service providers flexibil-

ity in choosing algorithms and vendors, reducing costs, and enabling services to be distributed across cloud, edge, and proprietary infrastructure to avoid vendor lock-in. This approach facilitates continuous service optimization based on user satisfaction, running costs, and energy consumption. To ensure seamless cross-cluster communications in multi-cluster deployments, attention is needed for both cluster management and connectivity between clusters.

Unfortunately, contemporary versions of Kubernetes are unable to facilitate multi-cluster deployments. This limitation extends to two fronts, the first one of which is the orchestration and management of the multiple clusters that facilitate the various XR services. The second one is the communication among these XR services, across cluster boundaries. This endeavour becomes quite challenging when considering that XR services, on top of the aforementioned QoS requirements, are characterized by functional requirements, such as the need for User Datagram Protocol (UDP) support that serves as the cornerstone for numerous streaming use-cases, which are of paramount importance in the frame of XR services [4]. Furthermore, edge clusters are usually limited in terms of computational resources, and thus any attempt at implementing cross-cluster networking in a manner that spans across the cloud-edge fabric should take resource consumption into consideration (especially in the case of resource-intensive services, such as XR). Finally, to fully harvest the benefits of multi-cluster deployments, an ideal solution should facilitate dynamic, multi-ownership deployment scenarios and have a singular cross-cluster control plane to optimally schedule the various workloads. Thus, the ideal solution for supporting cross-cluster networking for XR services should be capable of fulfilling the following requirements:

- to not significantly increase resource consumption,
- to not contribute towards significantly increasing end-to-end latency,
- to support UDP,
- to have a singular cross-cluster control plane for optimal workload scheduling, and
- to facilitate dynamic, multi-ownership deployment scenarios.

Towards achieving this goal, this paper is dedicated to examining numerous state-of-the-art multi-cluster connectivity & management frameworks, and to proposing a novel solution that caters to the intricacies of XR services. More specifically, the proposed solution consists of two contemporary frameworks whose modus operandi is based on the support of multi-cluster deployments. These frameworks are *Cluster API* for deploying and managing multiple clusters, and *Liqo* for cross-cluster internetworking.

2 MULTI-CLUSTER MANAGEMENT

Multi-cluster management is a critical aspect of modern computing infrastructure, enabling organizations to efficiently oversee and coordinate multiple clusters of resources.

KubeFed [5] is a multi-cluster management framework that enables each cluster to leverage its own local master.

Unfortunately, this approach comes at the cost of additional software resources to support localised autonomy and synchronization across the various clusters. Furthermore, KubeFed supports multi-vendor environments, but the level of support may vary depending on the specific vendor and their level of compatibility with the Kubernetes API.

Karmada [6], on the other hand, is based on a different approach that utilizes a custom API Server. This server operates as a centralized control point and resembles the standard Kubernetes API, while high-level resources like Deployments are handled by custom controllers instead of following the standard workflow. This is determined by policy constraints set through Custom Resource Definitions. However, this approach does not fully adhere to Kubernetes, so administrators cannot manage lower-level objects such as pods in the context of effective monitoring and debugging.

Terraform [7], a cloud-agnostic infrastructure provisioning tool, allows the creation of resources from various cloud services using a unified infrastructure-as-code approach. Despite its declarative approach for defining the desired end-state of infrastructure, the use of a single end-state file can lead to performance issues, particularly in multi-cluster deployments outside the same network. Additionally, unlike the other explored multi-cluster management tools, Terraform is not free, which could be a deterrent for smaller application developer groups.

Cluster API [8] is a free and open-source framework that brings Kubernetes-style APIs and support for the lifecycle management of (workload) Kubernetes clusters. Cluster API is not limited to a specific infrastructure vendor but is designed to accommodate different cloud providers. This enables consistency, portability, automation and repeatability in cluster deployments, ultimately widening the possibility of orchestrating heterogeneous and multi-cloud domains in a unified and vendor-neutral fashion.

3 MULTI-CLUSTER INTERCONNECTIVITY

While the aforementioned frameworks are capable of establishing multi-cluster management, they do not provide any form of multi-cluster interconnectivity functionality. Thankfully, various cloud-native solutions have emerged to provide such functionalities. In the frame of this work, we propose a taxonomy of tools enabling interconnectivity across *Kubernetes* clusters. This taxonomy mainly focuses on two categories: *Service Mesh* based approaches with multi-cluster support and *Overlay Networks*. The range of enabling tools is selected based on a combination of their perceived maturity and documentation quality, the empirical experience of the authors of this article, and their availability in open-source version. All of the following frameworks are capable of achieving multi-cluster connectivity. However, the ideal solution should be capable of fulfilling all of the aforementioned QoS, and functional requirements that are associated with XR services.

3.1 Service Mesh Solutions

A service mesh is a dedicated network infrastructure layer that manages communication between services in an application. It manages service requests, providing features like

service discovery, load balancing, encryption, and failover. A service mesh typically relies on proxies (so-called sidecars) to form a mesh at the component level or, for instance, the host level.

Istio [9], **Linkerd** [10] and **Consul** [11] are service mesh platforms designed for microservices integration, traffic management, policy enforcement, and telemetry data aggregation. They employ Sidecar proxies to handle traffic between services within a cluster, forming a microservice mesh offering service discovery, Layer 7 routing, circuit breakers, policy enforcement, and telemetry recording. However, a notable limitation is that, operating at Layer 7, they lack support for UDP traffic, a critical aspect for XR services that heavily rely on UDP for operations like video streaming. While a potential solution, such as CONNECT UDP, has been proposed, there is currently no practical option for supporting UDP communication in Layer 7 service meshes. Additionally, service mesh solutions with sidecar proxies introduce high application and latency overhead, which can be prohibitive for latency-sensitive applications.

The aforementioned service mesh solutions are capable of tackling challenges related to Layer 7 networking, but are insufficient to facilitate use cases requiring Layer 2/3 networking. On the contrary, **Network Service Mesh** (NSM) [12] offers a cloud-native network solution for microservices-based applications, emphasizing programmable and intelligent networking for Layer 2/3 connectivity. However, similar to other service mesh solutions, NSM lacks a cross-cluster control plane, preventing dynamic workload scheduling across diverse locations for optimal performance.

3.2 Overlay Network Solutions

An overlay network is a virtual (or logical) network that is established on top of an existing physical network. All nodes in an overlay network are connected to one another using virtual (or logical) links. Each link corresponds to a specified path in the underlying network topology. Overlay Network approaches use peering strategies (e.g., using VPN tunnels) to automatically interconnect clusters (and services). The core idea behind overlay networks is the facilitation of communication among microservices regardless of their location in a multi-cloud environment, through the overlay network. Whether for federation or edge-cloud scenarios, these dynamic network topologies can also be helpful as a building block that does not enforce a specific orchestration platform, but allows seamless communication across multiple sites.

Submariner [13] is a multi-cluster networking solution designed for cloud-native applications, offering seamless connectivity between Kubernetes clusters across various locations. It provides key functionalities such as Cross-cluster Layer 3 connectivity, service discovery, and network policy enforcement, enabling communication between services in different clusters, even across different cloud providers or data centers. While Submariner excels at achieving cross-cluster connectivity, it does not address workload orchestration and observability, leaving these aspects to static approaches or external tools.

Skupper [14] utilizes Virtual Application Networks (VANs) to address multi-cluster communication challenges,

creating virtual networks connecting applications and services in a hybrid cloud at Layer 7. In Kubernetes, Skupper forms a network with each namespace having a Skupper instance, constantly sharing information about exposed services to create awareness across instances. Through annotation, Skupper exposes services, creating proxy endpoints available in all network namespaces, and, like many service mesh solutions, operates on Layer 7 networking, lacking support for UDP.

Liqo [15] is a sophisticated open-source framework designed to enable seamless connectivity among clusters distributed across various geographical locations, encompassing on-premises environments, edge devices, and cloud infrastructures. Operating on a peer-to-peer model, the interconnection between peered clusters is implemented through secure VPN tunnels, which are dynamically established at the end of the peering process. The architecture leverages a virtualization approach where remote clusters are abstracted as virtual nodes within the local cluster. This abstraction allows for transparent communication between interconnected clusters, regardless of the underlying Container Network Interface (CNI) plugin. In the context of bidirectional peering, Liqo creates virtual nodes in each cluster, serving as representations of the resources provided by the remote cluster.

Furthermore, Liqo introduces the concept of offloading, enabling the reflection and the execution of workloads such as namespaces, services, and pods on these virtual nodes. This capability facilitates the exposure of services and execution of workloads in remote clusters. For instance, when a namespace is offloaded, Liqo dynamically creates a twin namespace in the remote cluster. This twin namespace allows pods and services to run seamlessly in a shared, cross-cluster environment. In the pod offloading scenario, the actual execution of pods and associated services is moved to a peered cluster, demonstrating the flexibility to optimize resource usage across clusters. This is particularly useful for demanding computing tasks like video processing, enabling efficient workload distribution based on the capabilities of different clusters. On the other hand, service offloading involves exposing only Kubernetes services on a remote cluster while retaining pod execution in the original cluster. This strategy allows for more selective offloading, giving flexibility in optimizing specific components of the application.

4 PROPOSED SOLUTION

Out of all the multi-cluster connectivity solutions that were explored in the previous section, only Liqo manages to satisfy all of the functional requirements that are associated with XR services in terms of providing **support for UDP**, establishing a **singular cross-cluster control plane** for optimal workload scheduling, and facilitating **dynamic, multi-ownership deployment scenarios**. Thus, Liqo has been selected as the backbone of the proposed solution. The proposed solution for establishing cross-cluster networking to support XR services consists of two contemporary frameworks that are capable of supporting multi-cluster deployments. These frameworks are Cluster API for deploying and managing multiple clusters, and Liqo for cross-

cluster internetworking. To the best of our knowledge, this combination is novel and unexplored from both the design and implementation viewpoints.

An advanced implementation of the proposed solution that is based on Observe-Orient-Decide-Act (OODA) loops is depicted in Fig.2. The management cluster is responsible for instantiating the infrastructure and deploying applications on top of that infrastructure. It ensures observability by offering a monitoring framework with monitoring agents natively deployed alongside the infrastructure and applications. The monitoring framework utilizes Prometheus¹ and Thanos², both of which natively support Kubernetes clusters. It features a distributed architecture where each cluster has a monitoring server deployed to scrape data from local infrastructure nodes and applications. The management cluster aggregates data from these monitoring servers. The management cluster is divided into multiple levels, including an Observe and Orient level that gathers data and deploys algorithms for system observability. This advanced monitoring framework supports machine learning techniques for producing enhanced data analysis, and alerts. Once alerts and insights are generated, they are consumed by the Decide level. This level employs intelligent algorithms to find the best decisions in the frame of processes such as scheduling service components onto the right node of the right cluster, deciding when and how to perform service migrations, and performing resource scaling, among others. Once the corresponding decision is made, it is forwarded to the Act level, where it is enforced, leveraging Liqo and Cluster API. In order to isolate services and ensure connectivity between components of the same service, the same namespace is deployed across all clusters that run at least one component of the service. The components are deployed inside the service's namespace, and the connectivity between the components across clusters is ensured by Liqo tunnels. For instance, the traffic sent to the star service by component D will enter the blue VPN tunnel and be forwarded to component B.

As previously discussed, aside from the aforementioned functional requirements that only the proposed solution can guarantee, there are two additional QoS requirements that are essential for supporting multi-cluster networking for XR services. Thus, the next section is dedicated to showcasing that the proposed solution does not produce a significant overhead in terms of **resource consumption** and **latency**.

5 EXPERIMENTAL EVALUATION & DISCUSSION

To showcase the validity of the proposed solution in terms of facilitating multi-cluster management & networking, two experiments were conducted. The first study focused on assessing the speed at which various cluster sizes and distributions could be automated and provisioned using the Cluster API. Thus, it examines the effect that the use of Cluster API has on the overall latency from the perspective of provisioning times. The second study investigated the efficiency of Liqo in the context of video streaming, an important functionality in XR services, in terms of end-to-end latency and resource consumption.

1. <https://prometheus.io/>

2. <https://thanos.io/>

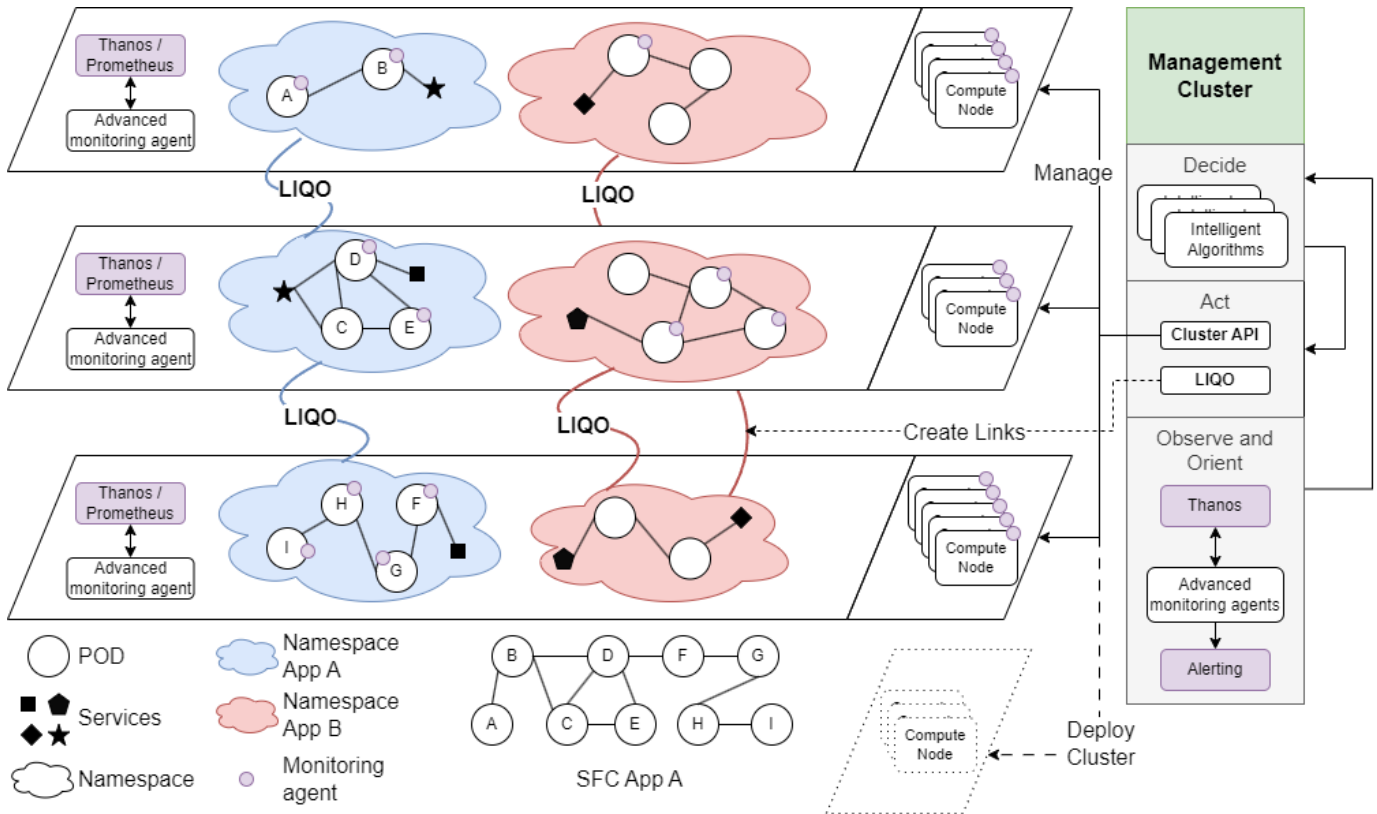


Fig. 2. An overview of the proposed solution.

5.1 Automated and Declarative Cluster Orchestration

Automated provisioning through Cluster API is crucial for minimizing infrastructure and tooling bootstrapping time, especially in impractical scenarios for larger Kubernetes clusters. This evaluation centers on automating and assessing the provisioning times for various Kubernetes cluster sizes and distributions, including lightweight options like K3s, prevalent in resource-constrained environments. The objective is to scrutinize Cluster API's support for diverse Kubernetes distributions and analyze the provisioning times for each.

In the frame of this experiment, infrastructure is conceptually divided into management and cloud, thus reflecting the two essential components of the proposed solution. The first consists of a Kubernetes cluster hosting the *Cluster API* (v1.3.1) components and a set of *Cluster API providers*, namely two Control Plane providers, *kubeadm* (v1.3.1) and *k3s* (v0.1.5) and one infrastructure provider, *OpenStack* (v0.7.0). Such management represents the key elements for orchestrating the remaining edge cloud. *Kubeadm* and *k3s* portray two widely used control plane installation options. On the other hand, the second part of the scenario consists of the cloud infrastructure using an OpenStack (Microstack Ussuri version) default installation to host the various Kubernetes clusters and their nodes where XR services will run.

The evaluation involved measuring the time required for generating cluster resource definitions (using *Cluster API syntax*), applying these resources to the management cluster, creating corresponding cloud resources (i.e., Virtual

Machines in OpenStack), and configuring the Kubernetes cluster. *Clusterctl*, part of *Cluster API*, was used to generate cluster definitions, leveraging default templates from each provider. For K3s on OpenStack, where no template existed, we created one³ specifically for Kubernetes clusters with *k3s* on OpenStack. The evaluation considered 1, 3, 5, 10, 15, 20, and 25 nodes for each Kubernetes distribution, representing one control plane, one control plane and two workers, and one control plane and four workers, respectively. Each node and Virtual Machine utilized *Ubuntu* images, with a flavour featuring 2vCPUs, 2GB RAM, and 20GB disk for both the control plane and worker nodes.

Fig. 3 displays the experiment results, depicting the total time for deploying all nodes. Deployment time is defined as the period until the control plane indicates *Ready* for single-node clusters or until all nodes indicate *Ready* for multi-node clusters. The tests were conducted 10 times, with error bars indicating the standard deviation. The results indicate that there is no significant difference in deployment time between the two Kubernetes distributions. Total deployment times ranged from 103 to 316 seconds, encompassing the creation of corresponding Virtual Machines on the target infrastructure. It is crucial to consider various factors influencing these times, including internal Cluster API reconciliation logic, cluster provisioning steps, target infrastructure, additional software installation (e.g., CNI), and checks required before marking a cluster as *Ready*. Additionally, scaling the number of nodes did not proportionally increase the total

3. <https://github.com/cluster-api-provider-k3s/cluster-api-k3s/pull/24>

deployment time, suggesting simultaneous bootstrapping of the control plane and remaining nodes by Cluster API. Thus, it is safe to conclude that the proposed solution is scalable for large cluster sizes.

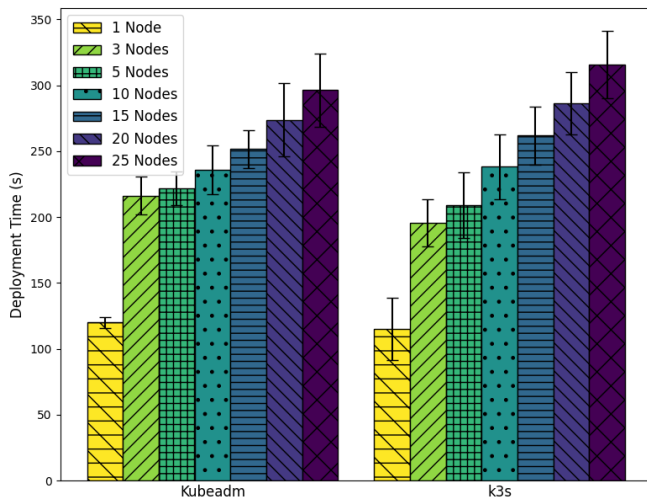


Fig. 3. Cluster API deployment time by Kubernetes distribution and cluster size.

5.2 Dynamic Cross-Cluster Networking

Since video streaming is a relevant aspect of XR services, we chose to evaluate the integration of Ligo to support cross-cluster video streaming. This evaluation aims to shed light on the effect that the use of Ligo has in terms of latency (due to overhead) and resource consumption. Towards achieving this goal, Ligo was compared against Kubernetes' NodePort. Although NodePort's approach of publicly exposing all services that need to communicate with each other prevents it from being a viable option for cross-cluster networking due to security and scalability concerns, it is capable of serving as a benchmark against Ligo in terms of end-to-end latency and resource consumption.

In the frame of the examined cross-cluster video streaming use-case, these nodes are used to host three key components. A Rendering Service (using `ffmpeg`⁴) sends a source feed using Hypertext Transfer Protocol (HTTP) to a Streaming Service. The Streaming Service (using `ffserver`⁵ as a media server) and a Client (another `ffmpeg` instance using Real Time Streaming Protocol (RTSP)) to consume the video. Based on this use-case, five scenarios were devised: 1) all components are deployed in the same Kubernetes cluster; 2) two local clusters on top of a singular OpenStack infrastructure connected via Ligo. In this scenario, one cluster hosts the rendering and client components, while the other accommodates the streaming service; 3) the same two local clusters of **sce.2**, but the services are exposed through NodePort; 4) two remote clusters, each one running on top of a different OpenStack infrastructure. One OpenStack deployment in Frankfurt and a second one in Geneva, connected through Ligo; 5) the same remote clusters of **sce.4**, but the services are exposed through NodePort;

4. <https://ffmpeg.org/>

5. <https://trac.ffmpeg.org/wiki/ffserver>

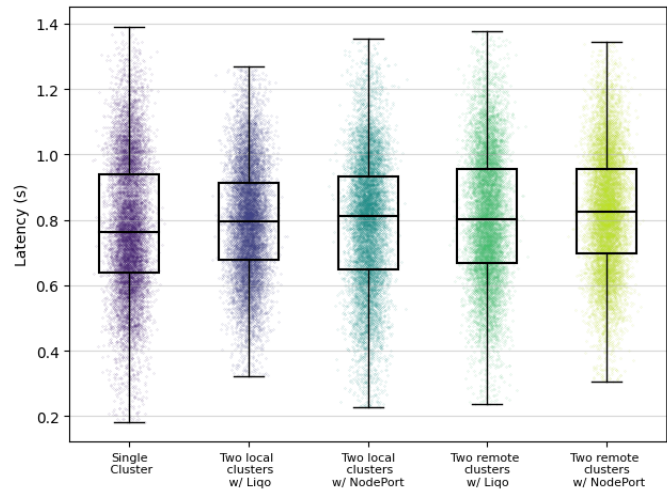


Fig. 4. Latency between streaming and view times for each scenario.

In all scenarios, each Kubernetes cluster was configured with 2 vCPUs, 4GB of RAM, and a 20GB disk. The setup utilized Ubuntu images for deployment. To assess Ligo overhead, we measure the overall end-to-end system latency (from the moment the frame is generated to the moment it is consumed). For each scenario, we conducted five independent runs of approximately 30 minutes of video streaming for a total of 150 min. For measuring latency, timestamps were embedded into the frames themselves using `ffmpeg` filters (generation and consumption timestamps, respectively), and later the difference was computed via Optical Character Recognition using Tesseract engine⁶. For each scenario, over 10000 frames were analysed. For each scenario, we also recorded the CPU and RAM consumption using Prometheus⁷. Fig. 4 shows the aggregated latency values of all runs per scenario. Whereas, Fig. 5 investigates resource consumption in the form of the CPU and Memory values. For scenarios with two clusters (**sce.2** to **sce.5**), the resource consumption values represent the average of the two.

Although **sce.1** is not aligned with the underlying premise of the next generation of XR services that dictates that these services shall be distributed across domains (and clusters), its results can provide us with hints about the overall latency regardless of the multi-domain aspect. According to the results plotted in Fig. 4, the 50th percentile for **sce.1** was equal to 762ms, thus surpassing the other scenarios. However, in all scenarios the 90th percentile consistently stayed around the 1-second mark, meaning that the significant majority of frames had an end-to-end latency equal to or below this value. Furthermore, **sce.1**, despite providing lower average values, presented slightly more data dispersion (with a standard deviation of 202ms) characterized by the occurrence of clustering around the 1s value. These results can be explained by the fact that all the processing occurred in the same cluster. Even so, it's important to note that no bottleneck took place according to the resource consumption results depicted in Fig. 5. Ligo

6. <https://github.com/tesseract-ocr/tesseract>

7. <https://prometheus.io/>

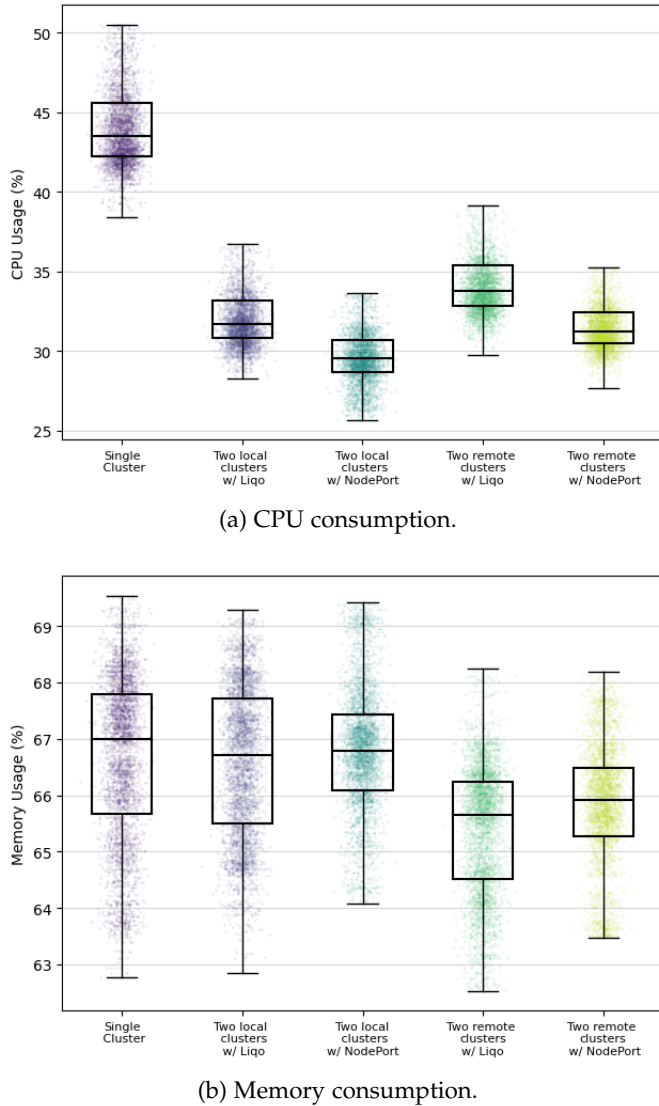


Fig. 5. Resource Consumption for each scenario.

performed slightly better compared to NodePort in the context of both local and remote clusters. More specifically, in **sce.2** and **sce.3**, the 50th percentiles were equal to 794ms and 811ms, while in **sce.4** and **sce.5** the 50th percentiles were equal to 801ms and 824ms. The low difference between local and remote clusters can be attributed to the relatively small physical distance among cluster locations (Geneva to Frankfurt). Although these values represent the overall end-to-end latency as perceived by a client in this particular scenario, and hence they are also dependent on the application itself, they demonstrate how Ligo’s overlay network is capable of supporting distributed XR architectures in terms of end-to-end latency.

Furthermore, the experimental results showcased in Fig. 5 indicate that the use of Ligo did not result in a prohibitively increased resource consumption in terms of CPU and memory. As expected, in **sce.1** the CPU usage was higher with a median of 43.2%. For the remaining scenarios, Ligo revealed relatively higher CPU consumption when compared to NodePort. More specifically, in **sce.2** and

sce.3, the 50th percentiles were equal to 31.5% and 29.4%, while in **sce.4** and **sce.5** the 50th percentiles were equal to 33.5% and 31.0%. These values can be attributed to the overlay strategy of having dedicated network tunnels between the clusters. Furthermore, the memory usage results do not indicate that the use of Ligo was accompanied by an increase in memory consumption. Instead, the examined scenarios exhibited an overall relatively similar behaviour varying approximately between 62% and 70%, regardless of whether or not they incorporated Ligo.

These findings are important, especially in the frame of resource-constrained environments. In such environments, using Ligo for cross-cluster networking introduces moderate CPU overhead due to its overlay network strategy, which remains manageable. While Ligo uses slightly more CPU than Kubernetes’ NodePort, its memory consumption is similar, making it suitable for memory-limited edge deployments. The key advantage of Ligo is its low latency, which is crucial for latency-sensitive XR services, thus providing efficient cross-cluster communication. This makes Ligo a viable solution for enabling distributed XR architectures in resource-constrained settings without significantly compromising resource availability. The experimental results that were explored in the frame of this section demonstrate that the proposed solution is capable of supporting not only the functional requirements of XR services, but also the aforementioned QoS requirements in terms of latency and resource consumption.

6 CONCLUSION

This work identified the Quality of Service (QoS) and functional requirements of XR services in multi-cluster deployments. After examining numerous contemporary technologies, the authors proposed a solution using two contemporary frameworks: Cluster API for multi-cluster orchestration and Ligo for multi-cluster networking. The solution was evaluated in two experiments. The first one assessed automation and provisioning times for different cluster sizes using Cluster API. The second tested Ligo’s efficiency in cross-cluster video streaming, a key use-case for XR services. The results showed that the solution effectively meets the aforementioned QoS and functional requirements of XR applications in multi-cluster environments.

In order to fully harvest the benefits of the proposed solution, it is vital to facilitate the operational complexity that arises from the need to orchestrate and monitor multiple clusters seamlessly in an automated manner that is aligned with the architecture that is depicted in Fig. 2. Thus, future research directions that spawn from the proposed solution involve the facilitation of the operational complexity that is associated with multi-cluster connectivity, and management. In the frame of establishing automated orchestration, it is essential to develop novel intelligent algorithms for handling dynamic workload distribution, service component placement, service migration, resource scaling, and many more. In the frame of developing comprehensive monitoring tools, it is of paramount importance to integrate machine learning techniques into monitoring frameworks to enable predictive analytics & anomaly detection, and to develop tools that support monitoring across federated Kubernetes

clusters and hybrid cloud environments, addressing challenges related to interoperability between on-premises and cloud-based Kubernetes deployments. Tackling these operational challenges is essential to effectively realizing the full potential of multi-cluster deployments for XR services.

ACKNOWLEDGEMENTS

The research leading to these results received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101016509 (project CHARITY). It is also supported in part by the European Union's HE research and innovation program HORIZON-JUSNS-2023 under the 6G-Path project (Grant No. 101139172). The paper reflects only the authors' views. The Commission is not responsible for any use that may be made of the information it contains. The work of JaeSeung Song was supported by the Ministry of Science and ICT (MSIT) under the Information Technology Research Center (ITRC) support program and the Technology Innovation Program funded By the Ministry of Trade, Industry & Energy (MOTIE, Korea) under Grant IITP-2024-2021-0-01816 and RS-2022-00154678. Prof. Song and Prof. Taleb are co-corresponding authors of this work.

REFERENCES

- [1] B. Burns, J. Beda, K. Hightower, and L. Evenson, *Kubernetes: up and running*. "O'Reilly Media, Inc.", 2022.
- [2] K. Boos, D. Chu, and E. Cuervo, "Demo: Flashback: Immersive virtual reality on mobile devices via rendering memoization," in *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services Companion*, ser. MobiSys '16 Companion. New York, NY, USA: Association for Computing Machinery, 2016, p. 94. [Online]. Available: <https://doi.org/10.1145/2938559.2938583>
- [3] T. Taleb, A. Boudi, L. Rosa, L. Cordeiro, T. Theodoropoulos, K. Tserpes, P. Dazzi, A. Protopsaltis, and R. Li, "Towards supporting xr services: Architecture and enablers," *IEEE Internet of Things Journal*, 2022.
- [4] J.-M. Chung, "Xr and multimedia video technologies," in *Emerging Metaverse XR and Video Multimedia Technologies: Modern Streaming and Multimedia Systems and Applications*. Springer, 2022, pp. 183–228.
- [5] "Kubefed", <https://github.com/kubernetes-sigs/kubefed>.
- [6] "Karmada", <https://karmada.io/>.
- [7] "TerraForm", <https://terraform.io/>.
- [8] "ClusterApi", <https://cluster-api.sigs.k8s.io/>.
- [9] "Istio", <https://istio.io/>.
- [10] "Linkerd", <https://linkerd.io/>.
- [11] "Consul", <https://consul.io/>.
- [12] "Network Service Mesh", <https://networkservicemesh.io/>.
- [13] "Submariner", <https://submariner.io/>.
- [14] "Skupper", <https://skupper.io/>.
- [15] M. Iorio, F. Risso, A. Palesandro, L. Camiciotti, and A. Manzalini, "Computing without borders: The way towards liquid computing," *IEEE Transactions on Cloud Computing*, vol. 11, no. 3, pp. 2820–2838, 2023.

Theodoros Theodoropoulos received the Eng. Diploma degree from the School of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece. He is currently pursuing the Ph.D. degree with the Department of Informatics and Telematics, Harokopio University of Athens, Kallithea, Greece. He has been working as a Research Engineer with the Harokopio University of Athens for the last three years. During this time, he had the chance to work at several Research and Development projects and to author numerous scientific

publications. His main research interests include deep learning, graph neural networks, deep reinforcement learning, and cloud & edge computing.

Luis Rosa received the Ph.D. degree in information science and technology from the University of Coimbra, Coimbra, Portugal, in 2021. He has been involved as a researcher in various European research projects. He has more than ten publications in journals, conferences, and book chapters on those topics. His research interests include edge/cloud computing, networks, security, and critical infrastructure protection.

Abderrahmane Boudi received the M.Sc. and Ph.D. degrees from the Higher National School of Computer Science (ESI), Algiers, Algeria, in 2013 and 2020 respectively. He was a visiting student of MOSAIC Lab/AALTO in 2017 and he has worked there as a research assistant in 2020. Currently, he is a senior lecturer at ESI and a researcher at ICTFICIAL Oy working on several projects on XR, AI, and Digital Twinning. His research interests include computing in the Edge/Cloud continuum, intelligent feedback control systems, SDN, and 6G and beyond networks.

Tarik Zakaria Benmerar received his engineer degree in computer engineering from USTHB-Algiers, in 2010 and master's degree in networks and distributed systems from USTHB-Algiers, in 2011. He received his PhD degree in parallelism and cloud computing (SaaS) applied for cerebral connectivity using diffusion MRI, at USTHB-Algiers in 2019. He is currently a senior lecturer at USTHB, and a research engineer at ICTFICIAL, Oy. His current research interests include cloud computing/edge/IoT architecture and orchestration, WebRTC-based streaming architectures and web browser-based parallelism.

Antonios Makris is a Senior Researcher at the School of Electrical and Computer Engineering at the National Technical University of Athens (NTUA). He received his BSc Degree in Computer Science in 2013 and MSc Degree in Web Engineering in 2015, both from Harokopio University of Athens. In 2022, he received his PhD in the area of Distributed Systems from the same department. His main research interests include Distributed Computing, Edge and Cloud Computing, Big Data Management and Analysis, Machine/Deep learning, NoSQL Database Systems and Spatiotemporal and Trajectory Analysis. He has participated in numerous European Union (EU) and National funded projects.

Tarik Taleb is currently a Full Professor at Ruhr University Bochum, Germany. He was a Professor with the Center of Wireless Communications (CWC), University of Oulu, Oulu, Finland. He is the founder of ICTFICIAL Oy, and the founder and the Director of the MOSAIC Lab, Espoo, Finland. From October 2014 to December 2021, he was an Associate Professor with the School of Electrical Engineering, Aalto University, Espoo, Finland. Prior to that, he was working as a Senior Researcher and a 3GPP Standards Expert with NEC Europe Ltd., Heidelberg, Germany. Before joining NEC and till March 2009, he worked as an Assistant Professor with the Graduate School of Information Sciences, Tohoku University, in a lab fully funded by KDDI. From 2005 to 2006, he was a Research Fellow with the Intelligent Cosmos Research Institute, Sendai. He received the B.E. degree (with distinction) in information engineering and the M.Sc. and Ph.D. degrees in information sciences from Tohoku University, Sendai, Japan, in 2001, 2003, and 2005, respectively.

Luis Cordeiro received the M.Sc. degree in communications and telematics from the University of Coimbra, Coimbra, Portugal, in 2007. He is a CTO with OneSource, Lisbon, Portugal. He has been actively involved in more than ten European research projects since 2004, mostly in the fields of networking, cloud, and security. He has several publications in journals, book chapters, and conferences. He has an extensive background in the areas of data communications, security, infrastructure management, and virtualization.

Konstantinos Tserpes is an Assistant Professor at the School of Electrical and Computer Engineering at the National Technical University of Athens (NTUA). He earned his PhD in Distributed Systems from the same institution. His research interests are primarily centered around efficient computing and data analytics systems, with a particular emphasis on the resource usage prediction on various applications. This work falls within the broader domain of systems engineering, where he has contributed significantly. He has also been an active participant in numerous European Union (EU) and National funded projects, where he has held various key roles, including that of scientific or general coordinator.

Jaesung Song is a full professor in both the Department of Computer & Information Security and the Department of Convergence Engineering for Intelligent Drones at Sejong University, Seoul, South Korea. He holds the position of Technical Plenary Vice Chair of the oneM2M global IoT standards initiative. Prior to his current position, he worked for NEC Europe Ltd. and LG Electronics in various positions. He received a Ph.D. at Imperial College London in the Department of Computing, United Kingdom. He holds B.S. and M.S. degrees in computer science from Sogang University. His research interests span the areas of beyond 5G and 6G, AI/ML enabled network systems, software engineering, networked systems and security, with focus on the design and engineering of reliable and intelligent IoT/M2M platforms.