
Liqo

The Liqo Authors

Oct 25, 2022

FEATURES

1	What does it provide?	3
2	What to explore next?	5
3	Peering	7
3.1	Overview	7
3.2	Approaches	8
4	Offloading	11
4.1	Virtual kubelet	11
4.2	Virtual node	11
4.3	Namespace extension	12
4.4	Pod offloading	12
4.5	Resource reflection	13
5	Network Fabric	15
5.1	Network manager	15
5.2	Cross-cluster VPN tunnels	15
5.3	In-cluster overlay network	16
6	Storage Fabric	17
7	Requirements	19
7.1	Resources	19
7.2	Connectivity	19
8	Ligo CLI tool	23
8.1	Introduction	23
8.2	Install liqoctl with Homebrew	23
8.3	Install liqoctl manually	24
8.4	Install Kubectrl plugin with Krew	25
8.5	Install liqoctl from source	25
8.6	Enable shell autocompletion	25
9	Install	29
9.1	Install with liqoctl	29
9.2	Customization options	34
9.3	Install with Helm	35
9.4	Install development versions	36
9.5	Ligo and Calico	36

10 Uninstall	39
10.1 Purge CRDs	39
11 Requirements	41
12 Quick Start	43
12.1 Provision the playground	43
12.2 Install Ligo	44
12.3 Peer two clusters	45
12.4 Leverage remote resources	46
12.5 Play with a microservice application	48
12.6 Tear down the playground	50
13 Offloading with Policies	53
13.1 Provision the playground	53
13.2 Peer the clusters	54
13.3 Tune namespace offloading	55
13.4 Deploy applications	56
13.5 Tear down the playground	57
14 Offloading a Service	59
14.1 Provision the playground	59
14.2 Peer the clusters	60
14.3 Offload a service	60
14.4 Tear down the playground	62
15 Stateful Applications	63
15.1 Provision the playground	63
15.2 Peer the clusters	64
15.3 Deploy a stateful application	64
15.4 Consume the database	65
15.5 Tear down the playground	66
16 Global Ingress	69
16.1 Provision the playground	69
16.2 Peer the clusters	70
16.3 Deploy an application	71
16.4 Check application spreading	71
16.5 Check service reachability	74
16.6 Tear down the playground	75
17 Peer two Clusters	77
17.1 Overview	77
17.2 Out-of-band control plane	78
17.3 In-band control plane	79
18 Namespace Offloading	83
18.1 Overview	83
18.2 Offloading a namespace	83
18.3 Unoffloading a namespace	85
19 Resource Reflection	87
19.1 Pods offloading	87
19.2 Service exposition	89
19.3 Persistent storage	90

19.4	Configuration data	90
20	Stateful Applications	91
20.1	Liqo virtual storage class	91
20.2	Externally managed storage	93
21	Prometheus Metrics	95
21.1	Scraping metrics	95
21.2	Cross-cluster network metrics	95
22	Contributing to Liqo	97
22.1	Repository structure	97
22.2	Release notes generation	97
22.3	Local development	97
22.4	Automatic tests	98

Liqo is an open-source project that enables **dynamic and seamless Kubernetes multi-cluster topologies**, supporting heterogeneous on-premise, cloud and edge infrastructures.

WHAT DOES IT PROVIDE?

Peering

Automatic peer-to-peer establishment of **resource and service consumption relationships** between independent and heterogeneous clusters. No need to worry about complex VPN configurations and certification authorities: everything is transparently **self-negotiated** for you.

Offloading

Seamless **workloads offloading** to remote clusters, without requiring any modification to Kubernetes or the applications themselves. **Multi-cluster is made native and transparent**: collapse an entire remote cluster to a **virtual node** compliant with the standard Kubernetes approaches and tools.

Network Fabric

A transparent **network fabric**, enabling multi-cluster **pod-to-pod** and **pod-to-service** connectivity, regardless of the underlying configurations and CNI plugins. Natively **access the services** exported by remote clusters, and spread interconnected application components across multiple infrastructures, with all cross-cluster traffic flowing through **secured network tunnels**.

Storage Fabric

A native **storage fabric**, supporting the remote execution of **stateful workloads** according to the **data gravity** approach. Seamlessly extend standard (e.g., database) **high availability deployment techniques** to the multi-cluster scenarios, for **increased guarantees**. All without the complexity of managing multiple independent cluster and application replicas.

WHAT TO EXPLORE NEXT?

Features

New to Ligo? Would you like to know more? Here you can find an in-depth overview of what a **peering** is, how the *virtual node* abstraction enables **workload offloading**, as well as discover about the **network and storage fabric** subsystems, ensuring the seamless functioning of unmodified multi-cluster applications.

Peering · Offloading · Network Fabric · Storage Fabric

Installation

Ready to give Ligo a try? Learn about installation and connectivity **requirements**, discover how to download and install **liqctl**, the CLI tool to streamline the installation and management of Ligo, and explore the **customization options**, based on the target environment characteristics.

Requirements · Ligo CLI tool · Install · Uninstall

Examples

Would you like to quickly join the fray and experiment with Ligo? Set up your playground and check out the **getting started examples**, which will guide you through a scenario-driven tour of the most notable features of Ligo. Discover how to **offload** (a subset of) your workloads, **access services** provided by remote clusters, **expose** multi-cluster applications, and more.

Quick Start · Offloading with Policies · Offloading a Service · Stateful Applications · Global Ingress

Usage

Do you want to make a step further and discover all the Ligo configuration options? These guides get you covered! Find out how to **establish and configure a peering** between two clusters, as well as how to enable and customize **namespace offloading**. Explore the details about which and how native resources are **reflected** to remote clusters, and learn more about the support for **stateful applications**.

Peer two Clusters · Namespace Offloading · Resource Reflection · Stateful Applications Prometheus Metrics

PEERING

In Ligo, we define **peering** a unidirectional resource and service consumption relationship between two Kubernetes clusters, with one cluster (i.e., the consumer) granted the capability to offload tasks to a remote cluster (i.e., the provider), but not vice versa. In this case, we say that the consumer establishes an **outgoing peering** towards the provider, which in turn is subjected to an **incoming peering** from the consumer.

This configuration allows for maximum flexibility in asymmetric setups, while transparently supporting bidirectional peerings through their combination. Additionally, the same cluster can play the role of provider and consumer in multiple peerings.

3.1 Overview

Overall, the establishment of a peering relationship between two clusters involves four main tasks:

- **Authentication:** each cluster, once properly authenticated through pre-shared tokens, obtains a valid identity to interact with the other cluster (i.e., its Kubernetes API server). This identity, granted only limited permissions concerning Ligo-related resources, is then leveraged to negotiate the necessary parameters, as well as during the offloading process.
- **Parameters negotiation:** the two clusters exchange the set of parameters required to complete the peering establishment, including the amount of resources shared with the consumer cluster, the information concerning the setup of the network VPN tunnel, and more. The process is completely automatic and requires no user intervention.
- **Virtual node setup:** the consumer cluster creates a new **virtual node** abstracting the resources shared by the provider cluster. This transparently enables the task offloading process detailed in the [offloading section](#), and it is completely compliant with standard Kubernetes practice (i.e., it requires no API modifications for application deployment and exposition).
- **Network fabric setup:** the two clusters configure their **network fabric** and establish a secure cross-cluster VPN tunnel, according to the parameters negotiated in the previous phase (endpoints, security keys, address remappings, ...). Essentially, this enables pods hosted by the local cluster to seamlessly communicate with the pods offloaded to a remote cluster, regardless of the underlying CNI plugin and configuration. Additional details are presented in the [network fabric section](#).

3.2 Approaches

Liqo supports two non-mutually exclusive peering approaches (i.e., the same cluster can leverage a different approach for different remote clusters), respectively referred to as **out-of-band control plane** and **in-band control plane**. The following sections briefly overview the differences among them, outlining the respective trade-offs. Additional in-depth details about the networking requirements are presented in the *installation requirements section*, while the *usage section* describes the operational commands to establish both types of peering.

3.2.1 Out-of-band control plane

The standard peering approach is referred to as **out-of-band control plane**, since the **Liqo control plane traffic** (i.e., including both the initial authentication process and the communication with the remote Kubernetes API server) **flows outside the VPN tunnel** interconnecting the two clusters (still, TLS is used to ensure secure communications). Indeed, this tunnel is dynamically started in a later stage of the peering process, and it is leveraged only for cross-cluster pods traffic.

The single cross-cluster traffic flow required by this approach is schematized at a high level in the figure below (agnostic from how services are exposed, which is presented in the *dedicated installation requirements section*).

Overall, the out-of-band control plane approach:

- Supports clusters under the control of **different administrative domains**, as each party interacts only with its own cluster: the provider retrieves an authentication token that is subsequently shared with and leveraged by the consumer to start the peering process.
- Is characterized by **high dynamism**, as upon parameters modifications (e.g., concerning VPN setup) the negotiation process ensures synchronization between clusters and the peering automatically re-converges to a stable status.
- Requires each cluster to expose **three different endpoints** (i.e., the Liqo authentication service, the Liqo VPN endpoint and the Kubernetes API server), making them accessible from the pods running in the remote cluster.

3.2.2 In-band control plane

The alternative peering approach is referred to as **in-band control plane**, since the **Liqo control plane traffic flows inside the VPN tunnel** interconnecting the two clusters. In this case, the tunnel is statically established at the beginning of the peering process (i.e., part of the negotiation process is carried out directly by the Liqo CLI tool), and it is leveraged from that moment on for all inter-cluster traffic. The three different cross-cluster traffic flows required by this approach are schematized at a high level in figure below (agnostic from how services are exposed, which is presented in the *dedicated installation requirements section*).

Overall, the in-band control plane approach:

- Requires the administrator starting the peering process to have **access to both clusters** (although with limited permissions), as the network parameters negotiation is performed through the Liqo CLI tool (which interacts at the same time with both clusters). The remainder of the peering process, instead, is completed as usual, although the entire communication flows inside the VPN tunnel.
- **Statically configures** the cross-cluster **VPN tunnel** at peering establishment time, hence requiring manual intervention in case of configuration changes causing connectivity loss.

- **Relaxes the connectivity requirements**, as only the Liqo VPN endpoint needs to be reachable from the pods running in the remote cluster. Specifically, the Kubernetes API service is not required to be exposed outside the cluster.

OFFLOADING

Workload offloading is enabled by the **virtual node** abstraction. A virtual node is spawned at the end of the peering process in the local (i.e., consumer) cluster, and represents (and aggregates) the subset of resources shared by the remote cluster. This solution enables the **transparent extension** of the local cluster, with the new node (and its capabilities) seamlessly taken into account by the vanilla Kubernetes scheduler when selecting the best place for the workloads execution. At the same time, this approach is fully compliant with the **standard Kubernetes APIs**, hence allowing to interact with and inspect offloaded pods just as if they were executed locally.

4.1 Virtual kubelet

The virtual node abstraction is implemented by an extended version of the **Virtual Kubelet project**. A virtual kubelet replaces a traditional kubelet when the controlled entity is not a physical node. In the context of Ligo, it interacts with both the local and the remote clusters (i.e., the respective Kubernetes API servers) to:

1. Create the **virtual node resource** and reconcile its status with respect to the negotiated configuration.
2. **Offload the local pods** scheduled onto the corresponding (virtual) node to the remote cluster, while keeping their status aligned.
3. Propagate and synchronize the **accessory artifacts** (e.g., *Services*, *ConfigMaps*, *Secrets*, ...) required for proper execution of the offloaded workloads, a feature we call **resource reflection**.

For each remote cluster, a different instance of the Ligo virtual kubelet is started in the local cluster, ensuring isolation and segregating the different authentication tokens.

4.2 Virtual node

A **virtual node** summarizes and abstracts the **amount of resources** (e.g., CPU, memory, ...) shared by a given remote cluster. Specifically, the virtual kubelet automatically propagates the negotiated configuration into the *capacity* and *allocatable* entries of the node status.

Node conditions reflect the current status of the node, with periodic and configurable **healthiness checks** performed by the virtual kubelet to assess the reachability of the remote API server. This allows to mark the node as *not ready* in case of repeated failures, triggering the standard Kubernetes eviction strategies based on the configured *pod tolerations* (e.g., to enforce service continuity).

Finally, each virtual node includes a set of **characterizing labels** (e.g., geographical region, underlying provider, ...) suggested by the remote cluster. This enables the enforcement of **fine-grained scheduling policies** (e.g., through *affinity* constraints), in addition to playing a key role in the namespace extension process presented below.

4.3 Namespace extension

To enable seamless workload offloading, **Liqo extends Kubernetes namespaces** across the cluster boundaries. Specifically, once a given namespace is selected for offloading (see the *namespace offloading usage section* for the operational procedure), Liqo proceeds with the automatic creation of **twin namespaces** in the subset of selected remote clusters.

Remote namespaces host the actual **Pods offloaded** to the corresponding cluster, as well as the **additional resources** propagated by the resource reflection process. This behavior is presented in the figure below, which shows a given namespace existing in the local cluster and extended to a remote cluster. A group of pods is contained in the local namespace, while a subset (i.e., those faded-out) is scheduled onto the virtual node and offloaded to the remote namespace. Additionally, the resource reflection process propagated different resources existing in the local namespace (e.g., *Services, ConfigMaps, Secrets, ...*) in the remote one (represented faded-out), to ensure the correct execution of offloaded pods.

The Liqo namespace extension process features a high degree of customization, mainly enabling to:

- Select a **specific subset of the available remote clusters**, by means of standard selectors matching the label assigned to the virtual nodes.
- Constraint whether pods should be scheduled onto **physical nodes only, virtual nodes only, or both**. The extension of a namespace, forcing at the same time all pods to be scheduled locally, enables the consumption of local services from the remote cluster, as shown in the *service offloading example*.
- Configure whether the **remote namespace name** should match the local one (although possibly incurring in conflicts), or be automatically generated, such as to be unique.

4.4 Pod offloading

Once a **pod is scheduled onto a virtual node**, the corresponding Liqo virtual kubelet (indirectly) creates a **twin pod object** in the remote cluster for actual execution. Liqo supports the offloading of both **stateless** and **stateful** pods, the latter either relying on the provided *storage fabric* or leveraging externally managed solutions (e.g., persistent volumes provided by the cloud provider infrastructure).

Remote pod resiliency (hence, service continuity), even in case of temporary connectivity loss between the two control planes, is ensured through a **custom resource** (i.e., *ShadowPod*) wrapping the pod definition, and triggering a Liqo enforcement logic running in the remote cluster. This guarantees that the desired pod is always present, without requiring the intervention of the originating cluster.

The virtual kubelet takes care of the automatic propagation of **remote status changes** to the corresponding local pod (remapping the appropriate information), allowing for complete **observability** from the local cluster. Advanced operations, such as **metrics and logs retrieval**, as well as **interactive command execution** inside remote containers, are transparently supported, to comply with standard troubleshooting operations.

Additional details concerning how pods are propagated to remote clusters are provided in the *resource reflection usage section*.

4.5 Resource reflection

The **resource reflection** process is responsible for the propagation and synchronization of selected control plane information into remote clusters, to enable the seamless execution of offloaded pods. Liqo supports the reflection of the resources dealing with **service exposition** (i.e., *Ingresses*, *Services* and *EndpointSlices*), **persistent storage** (i.e., *PersistentVolumeClaims* and *PersistentVolumes*), as well as those storing **configuration data** (i.e., *ConfigMaps* and *Secrets*).

All resources of the above types that live in a **namespace selected for offloading** are automatically propagated into the corresponding twin namespaces created in the selected remote clusters. Specifically, the local copy of each resource is the source of trust leveraged to realign the content of the **shadow copy** reflected remotely. Appropriate **remapping** of certain information (e.g., *endpoint IPs*) is transparently performed by the virtual kubelet, accounting for conflicts and different configurations in different clusters.

You can refer to the [resource reflection usage section](#) for a detailed characterization of how the different resources are reflected into remote clusters.

NETWORK FABRIC

The **network fabric** is the Ligo subsystem transparently extending the Kubernetes network model across multiple independent clusters, such that **offloaded pods can communicate with each other** as if they were all executed locally.

In detail, the network fabric ensures that **all pods in a given cluster can communicate with all pods on all remote peered clusters**, either with or without NAT translation. The support for arbitrary clusters, with different parameters and components (e.g., CNI plugins), makes it impossible to guarantee **non-overlapping pod IP address ranges** (i.e., *PodCIDR*). Hence, possibly requiring **address translation mechanisms**, provided that NAT-less communication is preferred whenever address ranges are disjointed.

The figure below represents at a high level the network fabric established between two clusters, with its main components detailed in the following.

5.1 Network manager

The **network manager** (not shown in figure) represents the **control plane** of the Ligo network fabric. It is executed as a pod, and it is responsible for the **negotiation of the connection parameters** with each remote cluster during the peering process.

It features an **IP Address Management (IPAM) plugin**, which deals with possible **network conflicts** through the definition of high-level NAT rules (enforced by the data plane components). Additionally, it exposes an interface consumed by the reflection logic to handle **IP addresses remapping**. Specifically, this is leveraged to handle the *translation of pod IPs* (i.e., during the synchronization process from the remote to the local cluster), as well as during *EndpointSlices reflection* (i.e., propagated from the local to the remote cluster).

5.2 Cross-cluster VPN tunnels

The interconnection between peered clusters is implemented through **secure VPN tunnels**, made with *WireGuard*, which are dynamically established at the end of the peering process, based on the negotiated parameters.

Tunnels are set up by the **Ligo gateway**, a component of the network fabric that is executed as a *privileged* pod on one of the cluster nodes. Additionally, it appropriately populates the **routing table**, and configures, by leveraging *iptables*, the **NAT rules** requested to comply with address conflicts.

Although this component is executed in the *host network*, it relies on a **separate network namespace** and **policy routing** to ensure isolation and prevent conflicts with the existing Kubernetes CNI plugin. Moreover, **active/standby high-availability** is supported, to ensure minimum downtime in case the main replica is restarted.

5.3 In-cluster overlay network

The **overlay network** is leveraged to **forward all traffic** originating from local pods/nodes, and directed to a remote cluster, **to the gateway**, where it will enter the VPN tunnel. The same process occurs on the other side, with the traffic that exits from the VPN tunnel entering the overlay network to reach the node hosting the destination pod.

Liqo leverages a **VXLAN**-based setup, which is configured by a network fabric component executed on all physical nodes of the cluster (i.e., as a *DaemonSet*). Additionally, it is also responsible for the population of the appropriate **routing entries** to ensure correct traffic forwarding.

STORAGE FABRIC

The Ligo **storage fabric** subsystem enables the seamless offloading of **stateful workloads** to remote clusters. The solution is based on two main pillars:

- **Storage binding deferral** until its first consumer is scheduled onto a given cluster (either local or remote). This ensures that **new storage pools** are created in the exact location where their associated pods have just been scheduled onto for execution.
- **Data gravity**, entering in action in the subsequent scheduling processes, and involving a set of **automatic policies** to attract pods in the appropriate cluster. This guarantees that pods requesting **existing pools of storage** (e.g., following a restart) are scheduled onto the cluster physically hosting the corresponding data.

These approaches extend standard Kubernetes practice to multi-cluster scenarios, simplifying at the same time the configuration of **high availability** and **disaster recovery** scenarios. To this end, one relevant use-case is represented by database instances that need to be replicated and synchronized across different clusters, which is shown in the *stateful applications example*.

Under the hood, the Ligo storage fabric leverages a **virtual storage class**, which embeds the logic to **create the appropriate storage pools** in the different clusters. Whenever a new *PersistentVolumeClaim (PVC)* associated with the virtual storage class is created, and its consumer is bound to a (possibly virtual) node, the Ligo logic goes into action, based on the target node:

- If it is a **local node**, PVC operations are remapped to a second one, associated with the corresponding **real storage class**, to transparently **provision the requested volume**.
- In case of **virtual nodes**, the reflection logic is responsible for creating the **remote shadow PVC**, remapped to the negotiated storage class, and **synchronizing the PersistentVolume information**, to allow pod binding.

In both cases, **locality constraints** are automatically embedded within the resulting *PersistentVolumes (PVs)*, to make sure each pod is scheduled only onto the cluster where the associated storage pools are available.

Additional details about the **configuration of the Ligo storage fabric**, as well as concerning the possibility to **move storage pools** among clusters through the Ligo CLI tool, are presented in the *stateful applications usage section*.

Note

In addition to the provided storage class, Ligo supports the execution of pods leveraging cross-cluster storage managed by external solutions (e.g., persistent volumes provided by the cloud provider infrastructure).

REQUIREMENTS

This page presents an overview of the main requirements, both in terms of **resources** and **network connectivity**, to use Ligo and successfully establish peerings with remote clusters.

7.1 Resources

Ligo requires very **limited resources** (i.e., CPU, RAM, network bandwidth), making it suitable for both traditional **K8s** clusters and **resource constrained** clusters, e.g., the ones running K3s on a Raspberry Pi.

While the exact numbers depend on the **number of established peerings**, **number of offloaded pods** and on the **size of the cluster**, as a ballpark figure the entire Ligo control plane, executed on a two-nodes KinD cluster, peered with one remote cluster, and while offloading 100 pods, requires less than:

- 0.5 CPU cores (only during transient periods, while CPU consumption is practically negligible in all the other instants).
- 200 MB of RAM (this metric increases the more pods are offloaded to remote clusters).
- 5 Mbps of cross-cluster control plane traffic (only during transient periods). Data plane traffic, instead, depends on the applications and their actual placements across the clusters.

However, to be on the safe side, we suggest installing Ligo on a cluster that has **at least 2 CPUs and 2 GB of RAM**, which takes into account also the resources used by standard Kubernetes components.

An accurate analysis of the Ligo performance compared to vanilla Kubernetes, including the characterization of the resources consumed by Ligo, is presented in a [dedicated blog post](#).

7.2 Connectivity

Ligo supports two alternative peering approaches, each characterized by **different requirements in terms of network connectivity** (i.e., mutually reachable endpoints):

- *Out-of-band control plane peering*: it requires **three separated traffic flows** (hence, three exposed endpoints).
- *in-band control plane peering*: it requires a **single endpoint**, as all control plane traffic is tunneled inside the cross-cluster VPN.

More details available in the [peering section](#).

Note

The two peering approaches are **non-mutually exclusive**: a cluster can leverage different approaches toward different remote clusters, provided that the connectivity requirements are satisfied.

7.2.1 Out-of-band control plane peering

In order to successfully establish an out-of-band control plane peering with a remote cluster, the following three services need to be **reciprocally accessible** on both clusters (i.e., in terms of IP address/port):

- **Authentication service** (`liqo-auth`): the Liqo service used to authenticate incoming peering requests coming from other clusters.
- **Network gateway** (`liqo-gateway`): the Liqo service responsible for the setup of the cross-cluster VPN tunnels.
- **Kubernetes API server**: the standard Kubernetes API Server, that is used by the (remote) Liqo instance to create the resources required to start the peering process, and perform workload offloading.

Reciprocally accessible means that a first cluster must be able to connect to the `<IP/port>` of the above services exposed on the second cluster, and vice versa (i.e., from second cluster to the first); some exceptions that refer to the network gateway are detailed in the following of this page. This implies also that any network device (**NAT**, **firewall**, etc.) sitting in the path between the two clusters must be configured to **enable direct connectivity** toward the above services, as presented in the *network firewalls* section.

The tuple `<IP/port>` exported by the Liqo services (i.e., `liqo-auth`, `liqo-gateway`) depends on the Liqo configuration, chosen at installation time, which may depend on the physical setup of your cluster and the characteristics of your service. In particular:

- **Authentication Service**: when you install Liqo, you can choose to expose the authentication service through a *LoadBalancer* service, a *NodePort* service, or an *Ingress* (the last allows the service to be exposed as *ClusterIP*). This choice depends (1) on your necessities, (2) on the cluster configuration (e.g., a *NodePort* cannot be used if your nodes have private IP addresses, hence cannot be reached from the Internet), and (3) whether the above primitives (e.g., the *Ingress Controller*) are available in your cluster.
- **Network Gateway**: the same applies also for the network gateway, except that it cannot be exported through an *Ingress*. In fact, while the authentication service uses a standard HTTP/REST interface, the network gateway is the termination of a UDP-based network tunnel; hence only *LoadBalancer* and *NodePort* services are supported.

Note

Liqo supports scenarios in which, given two clusters, only one of the two **network gateways** is publicly reachable from the remote cluster (i.e., in terms of `<IP/port>` tuple), although communication must be allowed by possible firewalls sitting in the path.

By default, `liqoctl` exposes both the authentication service and the network gateway through a **dedicated LoadBalancer service**, falling back to a *NodePort* for simpler setups (i.e., KinD and K3s). However, more advanced configurations can be achieved by configuring the proper [Helm chart parameters](#), either directly or by customizing the installation process *through liqoctl*.

An overview of the overall connectivity requirements to establish out-of-band control plane peerings in Liqo is shown in the figure below.

Additional considerations

The choice of the way you expose Liqo services to remote clusters may not be trivial in some cases. Here, we list some additional notes you should consider in your choice:

- **NodePort service:** although a *NodePort* service can be used to expose the authentication service and the network gateway, often the IP addresses of the nodes are configured with private IP addresses, hence not being suitable for connections originated from the Internet. This happens rather often in production clusters, and on public clusters as well.
- **Ingress controller:** in case the authentication service is exposed through an *Ingress*, you should remember that, by default, the authentication service uses the TLS protocol. Hence, either you configure your *Ingress Controller* to connect to backend services with TLS as well, or you disable TLS on the authentication service.

Finally, in some cases clusters may reside behind a NAT. Liqo transparently supports scenarios with **one cluster behind NAT** and the other publicly reachable. Yet, in such situations, we suggest leveraging the in-band peering, as it simplifies the overall configuration.

7.2.2 In-band control plane peering

In order to successfully establish an in-band control plane peering with a remote cluster, you need only the **network gateways to be mutually reachable**, since all the Liqo control plane traffic is then configured to flow inside the VPN tunnel. All considerations presented above and referring to the exposition of the network gateway apply also in this case.

Given the connectivity requirements are a subset of the previous case, this solution is compatible with the configurations that enable the out-of-band peering approach. Additionally, it:

- Supports scenarios characterized by a **non publicly accessible Kubernetes API Server**.
- Allows to expose the authentication service as a *ClusterIP* service, reducing the number of services exposed externally.
- Enables setups with one cluster **behind NAT**, since the VPN tunnel can be established successfully even in case only one of the two network gateways is publicly reachable from the other cluster.

An overview of the overall connectivity requirements to establish in-band control plane peerings in Liqo is shown in the figure below.

Warning: Due to current limitations, the establishment of an in-band peering may not complete successfully in case the authentication service is exposed through an Ingress to which the TLS termination is delegated (i.e., when TLS is disabled on the authentication service).

7.2.3 Network firewalls

In some cases, especially on production setups, additional network limitations are present, such as firewalls that may impair network connectivity, which must be considered in order to enable Liqo peerings.

Depending on your configuration and the selected peering approach, you may have to configure existing firewalls to enable remote clusters to contact either the `liqo-gateway` only or all the three endpoints (i.e., `liqo-auth`, `liqo-gateway` and Kubernetes API server) that need to be publicly accessible in the peering phase.

To know the network parameters (i.e., <IP/port>) used by `liqo-auth` and `liqo-gateway`, you can use standard Kubernetes commands (e.g., `kubectl get services -n liqo`), while the <IP/port> tuple used by your Kubernetes API server is the one written in the `kubeconfig` file.

Remember that the Kubernetes API server and authentication service use the HTTPS protocol (over TCP); vice versa, the network gateway uses the [WireGuard](#) protocol over UDP.

LIQO CLI TOOL

8.1 Introduction

Liqoctl is the CLI tool to streamline the **installation** and **management** of Ligo. Specifically, it abstracts the creation and modification of Ligo defined custom resources, allowing to:

- **Install/uninstall** Ligo, wrapping the corresponding Helm commands and automatically retrieving the appropriate parameters based on the target cluster configuration.
- Establish and revoke **peering** relationships towards remote clusters.
- Enable and configure **workload offloading** on a per-namespace basis.
- Retrieve the **status** of Ligo, as well as of given peering relationships and offloading setups.

Warning: Make sure to always use the *liqoctl* version matching that of Ligo installed (or to be installed) in your cluster(s).

Note

liqoctl displays a *kubectl* compatible behavior concerning Kubernetes API access, hence supporting the KUBECONFIG environment variable, as well as all the standard flags, including `--kubeconfig` and `--context`. Moreover, subcommands interacting with two clusters (e.g., *liqoctl peer in-band*) feature a parallel set of flags concerning Kubernetes API access to the remote cluster, in the form `--remote-<flag>` (e.g., `--remote-kubeconfig`, `--remote-context`).

8.2 Install liqoctl with Homebrew

If you are using the [Homebrew](#) package manager, you can install *liqoctl* with Homebrew:

```
brew install liqoctl
```

When installed with Homebrew, autocompletion scripts are automatically configured and should work out of the box.

8.3 Install liqoctl manually

You can download and install *liqoctl* manually, following the appropriate instructions based on your operating system and architecture.

Linux

Download *liqoctl* and move it to a file location in your system PATH:

AMD64:

```
curl --fail -LS "https://github.com/liqotech/liqo/releases/download/v0.6.0/liqoctl-linux-  
↪amd64.tar.gz" | tar -xz  
sudo install -o root -g root -m 0755 liqoctl /usr/local/bin/liqoctl
```

ARM64:

```
curl --fail -LS "https://github.com/liqotech/liqo/releases/download/v0.6.0/liqoctl-linux-  
↪arm64.tar.gz" | tar -xz  
sudo install -o root -g root -m 0755 liqoctl /usr/local/bin/liqoctl
```

Note

Make sure `/usr/local/bin` is in your PATH environment variable.

MacOS

Download *liqoctl*, make it executable, and move it to a file location in your system PATH:

Intel:

```
curl --fail -LS "https://github.com/liqotech/liqo/releases/download/v0.6.0/liqoctl-  
↪darwin-amd64.tar.gz" | tar -xz  
chmod +x liqoctl  
sudo mv liqoctl /usr/local/bin/liqoctl
```

Apple Silicon:

```
curl --fail -LS "https://github.com/liqotech/liqo/releases/download/v0.6.0/liqoctl-  
↪darwin-arm64.tar.gz" | tar -xz  
chmod +x liqoctl  
sudo mv liqoctl /usr/local/bin/liqoctl
```

Note

Make sure `/usr/local/bin` is in your PATH environment variable.

Windows

Download the *liqoctl* binary:

```
curl --fail -LSO "https://github.com/liqotech/liqo/releases/download/v0.6.0/liqoctl-  
→windows-amd64"
```

And move it to a file location in your system PATH.

Alternatively, you can manually download *liqoctl* from the [Liqo releases](#) page on GitHub.

8.4 Install Kubectl plugin with Krew

You can install *liqoctl* as a kubectl plugin by using the [Krew](#) plugin manager:

```
kubectl krew install liqo
```

Then, all commands shall be invoked with `kubectl liqo` rather than `liqoctl`, although all functionalities remain the same.

Warning: While the kubectl plugin is supported, it is recommended to use *liqoctl* as this enables a better experience via tab auto-completion. *Install it with Homebrew* if available on your system or by *manually downloading the binary from GitHub*.

8.5 Install liqoctl from source

You can install *liqoctl* building it from source. To do so, clone the Liqo repository, build the *liqoctl* binary, and move it to a file location in your system PATH:

```
git clone https://github.com/liqotech/liqo.git  
cd liqo  
make ctl  
mv liqoctl /usr/local/bin/liqoctl
```

8.6 Enable shell autocompletion

liqoctl provides autocompletion support for Bash, Zsh, Fish, and PowerShell.

Bash

The *liqctl* completion script for Bash can be generated with the `liqctl completion bash` command. Sourcing the completion script in your shell enables *liqctl* autocompletion.

However, the completion script depends on *bash-completion*, which means that you have to install this software first. You can test if you have *bash-completion* already installed by running `type _init_completion`. If it returns an error, you can install it via your OS's package manager.

To load completions in your current shell session:

```
source <(liqctl completion bash)
```

To load completions for every new session, execute once:

```
source <(liqctl completion bash) >> ~/.bashrc
```

After reloading your shell, *liqctl* autocompletion should be working.

Zsh

The *liqctl* completion script for Zsh can be generated with the `liqctl completion zsh` command.

If shell completion is not already enabled in your environment, you will need to enable it. You can execute the following once:

```
echo "autoload -U compinit; compinit" >> ~/.zshrc
```

To load completions for each session, execute once:

```
liqctl completion zsh > "${fpath[1]}/_liqctl"
```

After reloading your shell, *liqctl* autocompletion should be working.

Fish

The *liqctl* completion script for Fish can be generated with the `liqctl completion fish` command.

To load completions in your current shell session:

```
liqctl completion fish | source
```

To load completions for every new session, execute once:

```
liqctl completion fish > ~/.config/fish/completions/liqctl.fish
```

After reloading your shell, *liqctl* autocompletion should be working.

PowerShell

The *liqctl* completion script for PowerShell can be generated with the `liqctl completion powershell` command.

To load completions in your current shell session:

```
liqctl completion powershell | Out-String | Invoke-Expression
```

To load completions for every new session, add the output of the above command to your PowerShell profile.

After reloading your shell, *liqctl* autocompletion should be working.

INSTALL

The deployment of all the Ligo components is managed through a **Helm chart**.

We strongly recommend **installing Ligo using *liqctl***, as it automatically handles the required customizations for each supported provider/distribution (e.g., AWS, EKS, GKE, Kubeadm, etc.). Under the hood, *liqctl* uses [Helm 3](#) to configure and install the Ligo Helm chart available on the official repository.

Alternatively, *liqctl* can also be configured to output a **pre-configured values file**, which can be further customized and used for a manual installation with Helm.

9.1 Install with *liqctl*

Below, you can find the basic information to install and configure Ligo, depending on the selected **Kubernetes distribution** and/or **cloud provider**. By default, *liqctl install* installs the latest stable version of Ligo, although it can be tuned through the `--version` flag.

The remainder of this page, then, presents **additional customization options** which apply to all setups, as well as advanced options.

Note

liqctl displays a *kubectl* compatible behavior concerning Kubernetes API access, hence supporting the `KUBECONFIG` environment variable, as well as all the standard flags, including `--kubeconfig` and `--context`. Ensure you selected the correct target cluster before issuing *liqctl* commands (as you would do with *kubectl*).

Kubeadm

Supported CNIs

Ligo supports Kubernetes clusters using the following CNIs: [Flannel](#), [Calico](#), [Canal](#), [Weave](#). Additionally, partial support is provided for [Cilium](#), although with the limitations listed below.

Warning: If you are installing Ligo on a cluster using the **Calico** CNI, you **MUST** read the [dedicated configuration section](#) to avoid unwanted misconfigurations.

Ligo + Cilium limitations

Liqo

Currently, Liqo supports the Cilium CNI only when *kube-proxy* is enabled. Additionally, known limitations concern the impossibility of accessing the backends of *NodePort* and *LoadBalancer* services hosted on remote clusters, from a local cluster using Cilium as CNI.

Installation

Liqo can be installed on a Kubeadm cluster through:

```
liqoctl install kubeadm
```

By default, the cluster is assigned an automatically generated name, then leveraged during the peering and offloading processes. Alternatively, you can manually specify a desired name with the `--cluster-name` flag.

OpenShift

Supported versions

Liqo was tested running on OpenShift Container Platform (OCP) 4.8.

Installation

Liqo can be installed on an OpenShift Container Platform (OCP) cluster through:

```
liqoctl install openshift
```

By default, the cluster is assigned an automatically generated name, then leveraged during the peering and offloading processes. Alternatively, you can manually specify a desired name with the `--cluster-name` flag.

AKS

Supported CNIs

Liqo supports AKS clusters using the following CNIs: [Azure AKS - Kubenet](#) and [Azure AKS - Azure CNI](#).

Configuration

To install Liqo on AKS, you should first log in using the az CLI (if not already done):

```
az login
```

Before continuing, you should export a few variables about the properties of your cluster:

```
# The resource group where the cluster is created
export AKS_RESOURCE_GROUP=resource-group
# The name of AKS cluster resource on Azure
export AKS_RESOURCE_NAME=cluster-name
# The name of the subscription associated with the AKS cluster
export AKS_SUBSCRIPTION_ID=subscription-name
```

Note

During the installation process, you need read-only permissions on the AKS cluster and on the Virtual Networks, if your cluster leverages the Azure CNI.

Installation

Liqo can be installed on an AKS cluster through:

```
liqctl install aks --resource-group-name "${AKS_RESOURCE_GROUP}" \
  --resource-name "${AKS_RESOURCE_NAME}" \
  --subscription-name "${AKS_SUBSCRIPTION_ID}"
```

By default, the cluster is assigned the same name as that specified through the `--resource-name` parameter. Alternatively, you can manually specify a different name with the `--cluster-name` *liqctl* flag.

Note

If you are running an [AKS private cluster](#), you may need to set the `--disable-endpoint-check` *liqctl* flag, since the API Server in your kubeconfig may be different from the one retrieved from the Azure APIs.

Additionally, since your API Server is not accessible from the public Internet, you shall leverage the *in-band peering approach* towards the clusters not attached to the same Azure Virtual Network.

EKS

Supported CNIs

Liqo supports EKS clusters using the default CNI: [AWS EKS - amazon-vpc-cni-k8s](#).

Configuration

Liqo leverages **AWS credentials to authenticate peered clusters**. Specifically, in addition to the read-only permissions used to configure the cluster installation (i.e., retrieve the appropriate parameters), Liqo uses AWS users to map peering access to EKS clusters.

To install Liqo on EKS, you should first log in using the `aws` CLI (if not already done). This is widely documented on the [official CLI documentation](#). In a nutshell, after having installed the CLI, you have to set up your identity:

```
aws configure
```

Before continuing, you should export a few variables about the properties of your cluster:

```
# The name of the target cluster
export EKS_CLUSTER_NAME=cluster-name
# The AWS region where the cluster is deployed
export EKS_CLUSTER_REGION=cluster-region
```

Then, you should retrieve the cluster's kubeconfig, if you have not already. You may use the following CLI command:

```
aws eks --region ${EKS_CLUSTER_REGION} update-kubeconfig --name ${EKS_CLUSTER_NAME}
```

Installation

Liqo can be installed on an EKS cluster through:

```
liqctl install eks --eks-cluster-region=${EKS_CLUSTER_REGION} \
  --eks-cluster-name=${EKS_CLUSTER_NAME}
```

By default, the cluster is assigned the same name as that specified through the `--eks-cluster-name` parameter. Alternatively, you can manually specify a different name with the `--cluster-name` *liqctl* flag.

GKE

Supported CNIs

Liqo supports GKE clusters using the default CNI: Google GKE - VPC-Native.

Warning: Liqo does not support GKE Autopilot Clusters.

Configuration

To install Liqo on GKE, you should create a service account for *liqoctl*, granting the read rights for the GKE clusters (you may reduce the scope to a specific cluster if you prefer).

First, let's start exporting a few variables about the properties of your cluster and the service account to create:

```
# The name of the service account used by liqoctl to interact with GCP
export GKE_SERVICE_ACCOUNT_ID=liqoctl
# The path where the GCP service account is stored
export GKE_SERVICE_ACCOUNT_PATH=~/.liqo/gcp_service_account

# The ID of the GCP project where your cluster was created
export GKE_PROJECT_ID=project-id
# The GCP zone where your GKE cluster is executed (if you are using zonal GKE clusters)
export GKE_CLUSTER_ZONE=europe-west1-b
# The GCP region where your GKE cluster is executed (if you are using regional GKE
↪clusters)
export GKE_CLUSTER_REGION=europe-west1
# The name of the GKE resource on GCP
export GKE_CLUSTER_ID=liqo-cluster
```

Second, you should create a GCP service account. This will represent the identity used by *liqoctl* to query the information required to properly configure Liqo on your cluster. The service account can be created using:

```
gcloud iam service-accounts create ${GKE_SERVICE_ACCOUNT_ID} \
  --project="${GKE_PROJECT_ID}" \
  --description="The identity used by liqoctl during the installation process" \
  --display-name="liqoctl"
```

Third, you should grant the service account the rights to inspect the cluster and the virtual networks parameters:

```
gcloud projects add-iam-policy-binding ${GKE_PROJECT_ID} \
  --member="serviceAccount:${GKE_SERVICE_ACCOUNT_ID}@${GKE_PROJECT_ID}.iam.
↪gserviceaccount.com" \
  --role="roles/container.clusterViewer"
gcloud projects add-iam-policy-binding ${GKE_PROJECT_ID} \
  --member="serviceAccount:${GKE_SERVICE_ACCOUNT_ID}@${GKE_PROJECT_ID}.iam.
↪gserviceaccount.com" \
  --role="roles/compute.networkViewer"
```

Fourth, you should create and download a set of valid service accounts keys, as presented by the official documentation. The keys will be used by *liqoctl* to authenticate to GCP:

```
gcloud iam service-accounts keys create ${GKE_SERVICE_ACCOUNT_PATH} \
  --iam-account=${GKE_SERVICE_ACCOUNT_ID}@${GKE_PROJECT_ID}.iam.gserviceaccount.com
```

Finally, you should retrieve the cluster's kubeconfig, if you have not already. You may use the following CLI command, in case of zonal GKE clusters:

```
gcloud container clusters get-credentials ${GKE_CLUSTER_ID} \
  --zone ${GKE_CLUSTER_ZONE} --project ${GKE_PROJECT_ID}
```

or, in case of regional GKE clusters:

```
gcloud container clusters get-credentials ${GKE_CLUSTER_ID} \
  --region ${GKE_CLUSTER_REGION} --project ${GKE_PROJECT_ID}
```

The retrieved kubeconfig will be added to the currently selected file (i.e., based on the KUBECONFIG environment variable, with fallback to the default path `~/.kube/config`) or created otherwise.

Installation

Liqo can be installed on a zonal GKE cluster through:

```
liqoctl install gke --project-id ${GKE_PROJECT_ID} \
  --cluster-id ${GKE_CLUSTER_ID} \
  --zone ${GKE_CLUSTER_ZONE} \
  --credentials-path ${GKE_SERVICE_ACCOUNT_PATH}
```

or, in case of regional GKE clusters:

```
liqoctl install gke --project-id ${GKE_PROJECT_ID} \
  --cluster-id ${GKE_CLUSTER_ID} \
  --region ${GKE_CLUSTER_REGION} \
  --credentials-path ${GKE_SERVICE_ACCOUNT_PATH}
```

By default, the cluster is assigned the same name as that assigned in GCP. Alternatively, you can manually specify a different name with the `--cluster-name` *liqoctl* flag.

K3s

Installation

Liqo can be installed on a K3s cluster through:

```
liqoctl install k3s
```

By default, the cluster is assigned an automatically generated name, then leveraged during the peering and offloading processes. Alternatively, you can manually specify a desired name with the `--cluster-name` flag.

KinD

Installation

Liqo can be installed on a KinD cluster through:

```
liqoctl install kind
```

By default, the cluster is assigned an automatically generated name, then leveraged during the peering and offloading processes. Alternatively, you can manually specify a desired name with the `--cluster-name` flag.

Other

Configuration

To install Liqo on alternative Kubernetes distributions, it is necessary to manually retrieve three main configuration parameters:

- **API Server URL:** the Kubernetes API Server URL (defaults to the one specified in the kubeconfig).
- **Pod CIDR:** the range of IP addresses used by the cluster for the pod network.
- **Service CIDR:** the range of IP addresses used by the cluster for service VIPs.

Installation

Once retrieved the above parameters, Liqo can be installed on a generic cluster through:

```
liqoctl install --api-server-url=<API-SERVER-URL> \  
--pod-cidr=<POD-CIDR> --service-cidr=<SERVICE-CIDR>
```

By default, the cluster is assigned an automatically generated name, then leveraged during the peering and offloading processes. Alternatively, you can manually specify a desired name with the `--cluster-name` flag.

9.2 Customization options

The following lists the main **customization parameters** exposed by the `liqoctl install` commands, along with a brief description. Additionally, **arbitrary parameters** available in the Helm *values* file (the full list is provided in the dedicated [repository page](#)) can be modified through the `--set` flag, which supports the standard Helm syntax.

9.2.1 Global

The main global flags, besides those concerning the installation of *development versions*, include:

- `--enable-ha`: whether to enable the support for **high-availability of the Liqo components**, starting two replicas (in an active/standby configuration) of the **gateway** to ensure no cross-cluster connectivity downtime in case one of the replicas is restarted, as well as of the **controller manager**, which embeds the Liqo control plane logic.
- `--enable-metrics`: enable **metrics** exposition through **prometheus**.
- `--timeout`: configures the timeout for the completion of the installation/upgrade process. Once expired, the process is aborted and Liqo is rolled back to the previous version.
- `--verbose`: whether to enable verbose logs, providing additional information concerning the installation/upgrade process (i.e., for troubleshooting purposes).

9.2.2 Control plane

The main control plane flags include:

- `--cluster-name`: configures a **name identifying the cluster** in Liqo. This name is propagated to remote clusters during the peering process, and used to identify the corresponding virtual nodes and the technical resources leveraged for the negotiation process. Additionally, it is leveraged as part of the suffix to ensure namespace names uniqueness during the offloading process. In case a cluster name is not specified, it is defaulted to that of the cluster in the cloud provider, if any, or it is automatically generated.

- `--cluster-labels`: a set of **labels** (i.e., key/value pairs) **identifying the cluster in Liqo** (e.g., geographical region, Kubernetes distribution, cloud provider, ...) and automatically propagated during the peering process to the corresponding virtual nodes. These labels can be used later to **restrict workload offloading to a subset of clusters**, as detailed in the *namespace offloading usage section*.
- `--sharing-percentage`: the maximum percentage of available **cluster resources** that could be shared with remote clusters.

9.2.3 Networking

The main networking flags include:

- `--reserved-subnets`: the list of **private CIDRs to be excluded** from the ones used by Liqo to remap remote clusters in case of address conflicts, as already in use (e.g., the subnet of the cluster nodes). The Pod CIDR and the Service CIDR shall not be manually specified, as automatically included in the reserved list.

9.3 Install with Helm

To install Liqo directly with Helm, it is possible to proceed as follows:

1. Add the Liqo Helm repository:

```
helm repo add liqo https://helm.liqo.io/
```

2. Update the local Helm repository cache:

```
helm repo update
```

3. Generate a pre-configured values file with *liqoctl*:

```
liqoctl install <provider> [flags] --only-output-values
```

The resulting *values* file is saved in the current directory, as `values.yaml`, or in the path specified through the `--dump-values-path` flag.

Note

This step is optional, but it relieves the user from the retrieval of the set of necessary parameters depending on the target provider/distribution. Alternatively, the upstream values file can be retrieved through:

```
helm show values liqo/liqo > values.yaml
```

4. Appropriately configure the *values* file. The full list of options is provided in the dedicated [repository page](#).
5. Install Liqo:

```
helm install liqo liqo/liqo --namespace liqo \
  --values <path-to-values-file> --create-namespace
```

9.4 Install development versions

In addition to released versions (including alpha and beta candidates), *liqoctl* provides the possibility to install **development versions** of Liqo. Development versions include:

- All commits merged into the master branch of Liqo.
- The commits of *pull requests* to the Liqo repository, whose images have been built through the appropriate bot command.

The installation of a development version of Liqo can be triggered specifying a **commit *SHA*** through the `--version` flag. In this case, *liqoctl* proceeds to **clone the repository** (either from the official repository, or from a fork configured through the `--repo-url` flag) at the given revision, and to leverage the Helm chart therein contained:

```
liqoctl install <provider> --version <commit-sha> --repo-url <forked-repo-url>
```

Alternatively, the Helm chart can be retrieved from a **local path**, as configured through the `--local-chart-path` flag:

```
liqoctl install <provider> --version <commit-sha> --local-chart-path <path-to-local-
↪ chart>
```

9.5 Liqo and Calico

Liqo adds several interfaces to the cluster nodes to handle cross-cluster traffic routing. Those interfaces are intended to not interfere with the normal CNI job.

However, by default, Calico scans all existing interfaces on a node to detect network configurations and establish the correct routes. To prevent misconfigurations, Calico shall then be configured to skip Liqo-managed interfaces during this process. This is required if Calico is configured in *BGP* mode, while not in case the *VPC native setup* is leveraged.

In Calico v3.17 and above, this can be performed by patching the *Installation CR*, adding the following:

```
apiVersion: operator.tigera.io/v1
kind: Installation
metadata:
  name: default
spec:
  calicoNetwork:
    nodeAddressAutodetectionV4:
      skipInterface: liqo.*
    ...
  ...
```

For Calico versions prior to 3.17, instead, you should modify the *calico-node DaemonSet*, adding the appropriate environment variable to the *calico-node* container.

```
apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: calico-node
  namespace: kube-system
spec:
  template:
```

(continues on next page)

(continued from previous page)

```
spec:
  containers:
  - name: calico-node
    env:
    - name: IP_AUTODETECTION_METHOD
      value: skip-interface=liqo.*
    ...
```


UNINSTALL

Liqo can be uninstalled by leveraging the dedicated *liqctl* command:

```
liqctl uninstall
```

Alternatively, the same operation can be performed directly with Helm:

```
helm uninstall liqo --namespace liqo
```

Note

Due to current limitations, the uninstallation process might hang in case peerings are still established, or namespaces are selected for offloading. To this end, *liqctl* performs a set of pre-checks and aborts the process in case any of the above is found, requesting the administrator to **unpeer all clusters and unoffload all namespaces** with the dedicated *liqctl* commands.

10.1 Purge CRDs

By default, the uninstallation process does not remove the Liqo CRDs and the system namespaces. These operations can be performed by adding the `--purge` flag:

```
liqctl uninstall --purge
```


REQUIREMENTS

Before starting the tutorials below, you should ensure the following software is installed on your system:

- **Docker**, the container runtime.
- **Kubectl**, the command-line tool for Kubernetes.
- **Helm**, the package manager for Kubernetes.
- **curl**, to interact with the tutorial applications through HTTP/HTTPS.
- **KinD**, the Kubernetes in Docker runtime.
- **liqctl** command-line tool to interact with Ligo.

The following tutorials were tested on Linux, macOS, and Windows (WSL2 and Docker Desktop).

Warning: To prevent issues with tutorials leveraging more than two clusters, on some systems you may need to increase the maximum number of *inotify* watches:

```
sudo sysctl fs.inotify.max_user_watches=52428899  
sudo sysctl fs.inotify.max_user_instances=2048
```


QUICK START

This tutorial aims at presenting how to install Liko and practicing with its most notable capabilities. You will learn how to create a *virtual cluster* by peering two Kubernetes clusters and how to deploy a simple application on it.

12.1 Provision the playground

First, check that you are compliant with the *requirements*.

Then, let's open a terminal on your machine and launch the following script, which creates a pair of clusters with KinD. Each cluster is made by two nodes (one for the control plane and one as a simple worker):

```
git clone https://github.com/liqotech/liqo.git
cd liqo
git checkout v0.6.0
cd examples/quick-start
./setup.sh
```

12.1.1 Explore the playground

You can inspect the deployed clusters by typing:

```
kind get clusters
```

You should see a couple of entries:

```
milan
rome
```

This means that two KinD clusters are deployed and running on your host.

Then, you can simply inspect the status of the clusters. To do so, you can export the KUBECONFIG variable to specify the identity file for *kubectl* and *liqoctl*, and then contact the cluster.

By default, the kubeconfigs of the two clusters are stored in the current directory (`./liqo_kubeconf_rome`, `./liqo_kubeconf_milan`). You can export the appropriate environment variables leveraged for the rest of the tutorial (i.e., KUBECONFIG and KUBECONFIG_MILAN), and referring to their location, through the following:

```
export KUBECONFIG="$PWD/liqo_kubeconf_rome"
export KUBECONFIG_MILAN="$PWD/liqo_kubeconf_milan"
```

Note

We suggest exporting the kubeconfig of the first cluster as default (i.e., KUBECONFIG), since it will be the entry point of the virtual cluster and you will mainly interact with it.

On the first cluster, you can get the available pods by merely typing:

```
kubectl get pods -A
```

Similarly, on the second cluster, you can observe the pods in execution:

```
kubectl get pods -A --kubeconfig "$KUBECONFIG_MILAN"
```

If the above commands return each an output similar to the following, your clusters are up and ready.

NAMESPACE	NAME	READY	STATUS	
↪ RESTARTS	AGE			↪
kube-system	coredns-558bd4d5db-9vdr9	1/1	Running	0 ↪
↪	3m58s			
kube-system	coredns-558bd4d5db-tzdxg	1/1	Running	0 ↪
↪	3m58s			
kube-system	etcd-rome-control-plane	1/1	Running	0 ↪
↪	4m10s			
kube-system	kindnet-fcspl	1/1	Running	0 ↪
↪	3m58s			
kube-system	kindnet-q6qkm	1/1	Running	0 ↪
↪	3m42s			
kube-system	kube-apiserver-rome-control-plane	1/1	Running	0 ↪
↪	4m10s			
kube-system	kube-controller-manager-rome-control-plane	1/1	Running	0 ↪
↪	4m11s			
kube-system	kube-proxy-2c9b1	1/1	Running	0 ↪
↪	3m42s			
kube-system	kube-proxy-7nngv	1/1	Running	0 ↪
↪	3m58s			
kube-system	kube-scheduler-rome-control-plane	1/1	Running	0 ↪
↪	4m11s			
local-path-storage	local-path-provisioner-85494db59d-sk55	1/1	Running	0 ↪
↪	3m58s			

12.2 Install Liqo

You will now install Liqo on both clusters, using the following characterizing names:

- **rome**: the *local* cluster, where you will deploy and control the applications.
- **milan**: the *remote* cluster, where part of your workloads will be offloaded to.

You can install Liqo on the *Rome* cluster by launching:

```
liqoctl install kind --cluster-name rome
```

This command will generate the suitable configuration for your KinD cluster and then install Liqo.

Similarly, you can install Liqo on the *Milan* cluster by launching:

```
liqoctl install kind --cluster-name milan --kubeconfig "$KUBECONFIG_MILAN"
```

On both clusters, you should see the following output:

```
INFO Kubernetes clients successfully initialized
INFO Installer initialized
INFO Cluster configuration correctly retrieved
INFO Installation parameters correctly generated
INFO All Set! You can now proceed establishing a peering (liqoctl peer --help for more
↪information)
```

And the Liqo pods should be up and running:

```
kubectl get pods -n liqo
```

NAME	READY	STATUS	RESTARTS	AGE
liqo-auth-74c795d84c-x2p6h	1/1	Running	0	2m8s
liqo-controller-manager-6c688c777f-4lv9d	1/1	Running	0	2m8s
liqo-crd-replicator-6c64df5457-bq4tv	1/1	Running	0	2m8s
liqo-discovery-546fd594dc-kg6db	1/1	Running	0	2m8s
liqo-gateway-78cf7bb86b-pkdpt	1/1	Running	0	2m8s
liqo-metric-agent-5667b979c7-snmkg	1/1	Running	0	2m8s
liqo-network-manager-5b5cdcfcf7-scvd9	1/1	Running	0	2m8s
liqo-proxy-6674dd7bbd-kr2ls	1/1	Running	0	2m8s
liqo-route-7wsrx	1/1	Running	0	2m8s
liqo-route-sz75m	1/1	Running	0	2m8s
liqo-webhook-7bb6dbf88d-2hvwX	1/1	Running	0	2m8s

12.3 Peer two clusters

Once Liqo is installed in your clusters, you can establish new *peerings*. In this example, since the two API Servers are mutually reachable, you will use the *out-of-band peering approach*.

First, get the *peer command* from the *Milan* cluster:

```
liqoctl generate peer-command --kubeconfig "$KUBECONFIG_MILAN"
```

Second, copy and paste the command in the *Rome* cluster:

```
liqoctl peer out-of-band milan --auth-url [redacted] --cluster-id [redacted] --auth-
↪token [redacted]
```

Now, the Liqo control plane in the *Rome* cluster will contact the provided authentication endpoint providing the token to the *Milan* cluster to get its Kubernetes identity.

You can check the peering status by running:

```
kubectl get foreignclusters
```

The output should look like the following, indicating that the cross-cluster network tunnel has been established, and an outgoing peering is currently active (i.e., the *Rome* cluster can offload workloads to the *Milan* one, but not vice versa):

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	AUTHENTICATION	
↔AGE						
milan	OutOfBand	Established	None	Established	Established	↔
↔12s						

At the same time, you should see a virtual node (`liqo-milan`) in addition to your physical nodes:

```
kubectl get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
liqo-milan	Ready	agent	14s	v1.23.6
rome-control-plane	Ready	control-plane,master	7m56s	v1.23.6
rome-worker	Ready	<none>	7m25s	v1.23.6

12.4 Leverage remote resources

Now, you can deploy a standard Kubernetes application in a multi-cluster environment as you would do in a single cluster scenario (i.e. no modification is required).

12.4.1 Start a hello world application

If you want to deploy an application that is scheduled onto Liqo virtual nodes, you should first create a namespace where your pod will be started. Then tell Liqo to make this namespace eligible for the pod offloading.

```
kubectl create namespace liqo-demo  
liqoctl offload namespace liqo-demo
```

The `liqoctl offload namespace` command enables Liqo to offload the namespace to the remote cluster. Since no further configuration is provided, Liqo will add a suffix to the namespace name to make it unique on the remote cluster (see the dedicated [usage page](#) for additional information concerning namespace offloading configurations).

Note

The virtual nodes have a `taint` that prevents the pods from being scheduled on them. The Liqo webhook will add the toleration for this taint to the pods created in the liqo-enabled namespaces.

Then, you can deploy a demo application in the `liqo-demo` namespace of the local cluster:

```
kubectl apply -f ./manifests/hello-world.yaml -n liqo-demo
```

The `hello-world.yaml` file represents a simple `nginx` service. It contains two pods running an `nginx` image and a Service exposing the pods to the cluster. One pod is running in the local cluster, while the other is forced to be scheduled on the remote cluster.

Info

Differently from the traditional examples, the above deployment introduces an *affinity* constraint. This forces Kubernetes to schedule the first pod (i.e. `nginx-local`) on a physical node and the second (i.e. `nginx-remote`) on a virtual node. Virtual nodes are like traditional Kubernetes nodes, but they represent remote clusters and have the `liqo.io/type: virtual-node` label.

When the affinity constraint is not specified, the Kubernetes scheduler selects the best hosting node based on the available resources. Hence, each pod can be scheduled either in the *local* cluster or in the *remote* cluster.

Now you can check the status of the pods. The output should be similar to the one below, confirming that one `nginx` pod is running locally; while the other is hosted by the virtual node (i.e., `liqo-milan`).

```
kubectl get pod -n liqo-demo -o wide
```

And the output should look like this:

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE	NOMINATED
↪NODE	READINESS GATES						
nginx-local	1/1	Running	0	94s	10.200.1.10	rome-worker	<none>
↪	<none>						
nginx-remote	1/1	Running	0	94s	10.202.1.9	liqo-milan	<none>
↪	<none>						

Check the pod connectivity

Once both pods are correctly running, it is possible to check one of the abstractions introduced by Liqo. Indeed, Liqo enables each pod to be transparently contacted by every other pod and physical node (according to the Kubernetes model), regardless of whether it is hosted by the *local* or by the *remote* cluster.

First, let's retrieve the IP address of the `nginx` pods:

```
LOCAL_POD_IP=$(kubectl get pod nginx-local -n liqo-demo --template={{.status.podIP}})
REMOTE_POD_IP=$(kubectl get pod nginx-remote -n liqo-demo --template={{.status.podIP}})
echo "Local Pod IP: ${LOCAL_POD_IP} - Remote Pod IP: ${REMOTE_POD_IP}"
```

You can fire up a pod and run `curl` from inside the cluster:

```
kubectl run --image=curlimages/curl curl -n default -it --rm --restart=Never -- curl $
↪{LOCAL_POD_IP}
kubectl run --image=curlimages/curl curl -n default -it --rm --restart=Never -- curl $
↪{REMOTE_POD_IP}
```

Both commands should lead to a successful outcome (i.e., return a demo web page), regardless of whether each pod is executed locally or remotely.

12.4.2 Expose the pods through a Service

The above `hello-world.yaml` manifest additionally creates a Service which is designed to serve traffic to the previously deployed pods. This is a traditional [Kubernetes Service](#) and can work with Liqo with no modifications.

Indeed, inspecting the Service, it is possible to observe that both `nginx` pods are correctly specified as endpoints. Nonetheless, it is worth noticing that the first endpoint (i.e. `10.200.1.10:80` in this example) refers to a pod running in the *local* cluster, while the second one (i.e. `10.202.1.9:80`) points to a pod hosted by the *remote* cluster.

```
kubectl describe service liqo-demo -n liqo-demo
```

```
Name:          liqo-demo
Namespace:    liqo-demo
```

(continues on next page)

Labels:	<none>
Annotations:	<none>
Selector:	app=liqo-demo
Type:	ClusterIP
IP Family Policy:	SingleStack
IP Families:	IPv4
IP:	10.93.84.150
IPs:	10.93.84.150
Port:	web 80/TCP
TargetPort:	web/TCP
Endpoints:	10.200.1.10:80,10.202.1.9:80
Session Affinity:	None
Events:	<none>

Check the Service connectivity

It is now possible to contact the Service: as usual, Kubernetes will forward the HTTP request to one of the available back-end pods. Additionally, all traditional mechanisms still work seamlessly (e.g. DNS discovery), even though one of the pods is actually running in a *remote* cluster.

You can fire up a pod and run `curl` from inside the cluster:

```
kubectrl run --image=curlimages/curl curl -n default -it --rm --restart=Never -- \
  curl --silent liqo-demo.liqo-demo.svc.cluster.local | grep 'Server'
```

Note

Executing the previous command multiple times, you will observe that part of the requests are answered by the pod running in the *local* cluster, and in part by that in the *remote* cluster (i.e., the `Server` value changes).

12.5 Play with a microservice application

It is very common in a cloud-based environment to deploy microservices applications composed of many pods interacting among each other. This pattern is transparently supported by Liqo and the virtual cluster abstraction.

You can play with a [microservices application](#) provided by Google, which includes multiple cooperating Services leveraging different networking protocols:

```
kubectrl apply -k ./manifests/demo-application -n liqo-demo
```

By default, Kubernetes schedules each pod either in the local or in the remote cluster, optimizing each deployment based on the available resources. However, you can play with *affinity* constraints to force Kubernetes to schedule of each component in a specific location, and see that everything continues to work smoothly. Specifically, the manifest above forces the frontend component to be executed in the *local* cluster, as this is required to enable *port-forwarding*, which is leveraged below.

Each demo component is exposed as a Service and accessed by other components. However, given that nobody knows, a priori, where each Service will be deployed (either locally or in the remote cluster), Liqo *replicates* all Kubernetes Services across both clusters, although the corresponding pod may be running only in one location. Hence, each microservice deployed across clusters can reach the others seamlessly: independently of the cluster a pod is deployed

in, each pod can contact other Services and leverage the traditional Kubernetes discovery mechanisms (e.g., DNS discovery and environment variables).

Additionally, several other objects (e.g. ConfigMaps and Secrets) inside a namespace are replicated in the remote cluster within the *twin namespace*, thus, ensuring that complex applications can work seamlessly across clusters.

12.5.1 Observe the application deployment

Once the demo application manifest is applied, you can observe the creation of the different pods:

```
watch kubectl get pods -n liqo-demo -o wide
```

At steady-state, you should see an output similar to the following. Different pods may be hosted by either the local nodes (*rome-worker* in the example below) or remote cluster (*liqo-milan* in the example below), depending on the scheduling decisions.

NAME	READY	STATUS	RESTARTS	AGE	IP
↔ NODE					
↔ NOMINATED NODE					
↔ READINESS GATES					
adservice-66f6b5c6fd-w95th	1/1	Running	0	2m23s	10.202.1.
↔19 liqo-milan <none>	<none>				
cartservice-76dc758684-wd5px	1/1	Running	0	2m23s	10.202.1.
↔15 liqo-milan <none>	<none>				
checkoutservice-85b74f746f-lm7gh	1/1	Running	0	2m24s	10.202.1.
↔11 liqo-milan <none>	<none>				
currencyservice-64775746dd-k85z6	1/1	Running	0	2m23s	10.202.1.
↔16 liqo-milan <none>	<none>				
emailservice-58f8b4f854-9mx7g	1/1	Running	0	2m24s	10.202.1.
↔10 liqo-milan <none>	<none>				
frontend-7b648dcb8f-gfbh2	1/1	Running	0	2m23s	10.200.1.
↔15 rome-worker <none>	<none>				
nginx-local	1/1	Running	0	6m5s	10.200.1.
↔10 rome-worker <none>	<none>				
nginx-remote	1/1	Running	0	6m5s	10.202.1.9
↔ liqo-milan <none>	<none>				
paymentservice-5dd7bb5855-ssbhf	1/1	Running	0	2m23s	10.202.1.
↔14 liqo-milan <none>	<none>				
productcatalogservice-587c8dbf7d-nmw77	1/1	Running	0	2m23s	10.202.1.
↔13 liqo-milan <none>	<none>				
recommendationservice-6cd468f4d4-h8k2t	1/1	Running	0	2m24s	10.202.1.
↔12 liqo-milan <none>	<none>				
redis-cart-78746d49dc-rkqrn	1/1	Running	0	2m23s	10.202.1.
↔18 liqo-milan <none>	<none>				
shippingservice-59c7b7458d-sqb9x	1/1	Running	0	2m23s	10.202.1.
↔17 liqo-milan <none>	<none>				

12.5.2 Access the demo application

Once the deployment is up and running, you can start using the demo application and verify that everything works correctly, even if its components are distributed across multiple Kubernetes clusters.

By default, the frontend web-page is exposed through a LoadBalancer Service, which can be inspected using:

```
kubectl get service -n liqo-demo frontend-external
```

Leverage `kubectl port-forward` to forward the requests from your local machine (i.e., `http://localhost:8080`) to the frontend Service:

```
kubectl port-forward -n liqo-demo service/frontend-external 8080:80
```

Open the `http://localhost:8080` page in your browser and enjoy the demo application.

12.6 Tear down the playground

12.6.1 Unoffload namespaces

Before starting the uninstallation process, make sure that all namespaces are unoffloaded:

```
liqoctl unoffload namespace liqo-demo
```

Every pod that was offloaded to a remote cluster is going to be rescheduled onto the local cluster.

12.6.2 Revoke peerings

Similarly, make sure that all the peerings are revoked:

```
liqoctl unpeer out-of-band milan
```

At the end of the process, the virtual node is removed from the local cluster.

12.6.3 Uninstall Liqo

Now you can uninstall Liqo from your clusters with `liqoctl`:

```
liqoctl uninstall  
liqoctl uninstall --kubeconfig="$KUBECONFIG_MILAN"
```

Purge

By default the Liqo CRDs will remain in the cluster, but they can be removed with the `--purge` flag:

```
liqoctl uninstall --purge  
liqoctl uninstall --purge --kubeconfig="$KUBECONFIG_MILAN"
```

12.6.4 Destroy clusters

To teardown the KinD clusters, you can issue:

```
kind delete cluster --name rome  
kind delete cluster --name milan
```


OFFLOADING WITH POLICIES

This tutorial aims to guide you through a tour to learn how to use the core Liko features. You will learn how to tune namespace offloading, and specify the target clusters through the *cluster selector* concept.

More specifically, you will configure a scenario composed of a *single entry point cluster* leveraged for the deployment of the applications (i.e., the *Venice* cluster, located in *north* Italy) and two *worker clusters* characterized by different geographical regions (i.e., the *Florence* and *Naples* clusters, respectively located in *center* and *south* Italy). Then, you will offload a given namespace (and the applications contained therein) to a subset of the worker clusters (i.e., only to the *Naples* cluster), while allowing pods to be also scheduled on the local cluster (i.e., the *Venice* one).

13.1 Provision the playground

First, check that you are compliant with the *requirements*.

Then, let's open a terminal on your machine and launch the following script, which creates the three above-mentioned clusters with KinD and installs Liko on all of them. Each cluster is made by a single combined control-plane + worker node.

```
git clone https://github.com/liqotech/liqo.git
cd liqo
git checkout v0.6.0
cd examples/offloading-with-policies
./setup.sh
```

Export the kubeconfigs environment variables to use them in the rest of the tutorial:

```
export KUBECONFIG="$PWD/liqo_kubeconf_venice"
export KUBECONFIG_FLORENCE="$PWD/liqo_kubeconf_florence"
export KUBECONFIG_NAPLES="$PWD/liqo_kubeconf_naples"
```

Note

We suggest exporting the kubeconfig of the first cluster as default (i.e., KUBECONFIG), since it will be the entry point of the virtual cluster and you will mainly interact with it.

At this point, you should have three clusters with Liko installed on them. The setup script named them **venice**, **florence** and **naples**, and respectively configured the following cluster labels:

- *venice*: topology.liqo.io/region=north
- *florence*: topology.liqo.io/region=center

- *naples*: topology.liqo.io/region=south

You can check that the clusters are correctly labeled through:

```
liqctl status
liqctl --kubeconfig $KUBECONFIG_FLORENCE status
liqctl --kubeconfig $KUBECONFIG_NAPLES status
```

These labels will be propagated to the virtual nodes corresponding to each cluster. In this way, you can easily identify the clusters through their characterizing labels, and define the appropriate scheduling policies.

13.2 Peer the clusters

Once Liqo is installed in your clusters, you can establish new *peerings*. In this example, since the two API Servers are mutually reachable, you will use the *out-of-band peering approach*.

To implement the desired scenario, let's first retrieve the *peer command* from the *Florence* and *Naples* clusters:

```
PEER_FLORENCE=$(liqctl generate peer-command --only-command --kubeconfig $KUBECONFIG_
↪FLORENCE)
PEER_NAPLES=$(liqctl generate peer-command --only-command --kubeconfig $KUBECONFIG_
↪NAPLES)
```

Then, establish the peerings from the *Venice* cluster:

```
echo "$PEER_FLORENCE" | bash
echo "$PEER_NAPLES" | bash
```

When the above commands return successfully, you can check the peering status by running:

```
kubectl get foreignclusters
```

The output should look like the following, indicating that an outgoing peering is currently active towards both the *Florence* and the *Naples* clusters, as well as the cross-cluster network tunnels have been established:

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	
↪AUTHENTICATION	AGE				
florence	OutOfBand	Established	None	Established	Established ↪
↪ 111s					
naples	OutOfBand	Established	None	Established	Established ↪
↪ 98s					

Additionally, you should have two new virtual nodes in the *Venice* cluster, characterized by the install-time provided labels:

```
kubectl get node --selector=liqo.io/type=virtual-node --show-labels
```

NAME	STATUS	ROLES	AGE	VERSION	LABELS
liqo-florence	Ready	agent	2m30s	v1.23.6	liqo.io/remote-cluster-id=4858caa2- ↪7848-4aab-a6b4-6984389bac56,liqo.io/type=virtual-node,topology.liqo.io/region=center
liqo-naples	Ready	agent	2m29s	v1.23.6	liqo.io/remote-cluster-id=53d11339- ↪fc85-4741-a912-a3eda781a69e,liqo.io/type=virtual-node,topology.liqo.io/region=south

Note

Some of the default labels were omitted for the sake of clarity.

13.3 Tune namespace offloading

Now, let's suppose you want to deploy an application that needs to be scheduled in the *north* and in the *south* region, but not in the *center* one. This constraint needs to be respected at the infrastructural level: the dev team does not need to be aware of required affinities and/or node selectors, nor it should be able to bypass them.

First, you should create a new namespace in the *Venice* cluster, which will host the application:

```
kubectl create namespace liqo-demo
```

Then, enable Liqo offloading for that namespace:

```
liqoctl offload namespace liqo-demo \
  --namespace-mapping-strategy EnforceSameName \
  --pod-offloading-strategy LocalAndRemote \
  --selector 'topology.liqo.io/region=south'
```

The above command configures the following aspects (see the dedicated [usage page](#) for additional information concerning namespace offloading configurations):

- the `liqo-demo` namespace is replicated with the same name in the other clusters.
- the `liqo-demo` namespace, and the contained resources, are offloaded only to the clusters with the `topology.liqo.io/region=south` label.
- the pods living in the `liqo-demo` namespace are free to be scheduled onto both physical and virtual nodes.

The `NamespaceOffloading` resource created by `liqoctl` in the `liqo-demo` namespace exposes the status of the offloading process, including a global `OffloadingPhase`, which is expected to be `Ready`, and a list of `RemoteNamespaceConditions`, one for each remote cluster.

In this case:

- the *Florence* cluster has not been selected to offload the namespace `liqo-demo`, since it does not match the cluster selector;
- the *Naples* cluster has been selected to offload the namespace `liqo-demo`, and the namespace has been correctly created.

```
kubectl get namespaceoffloadings offloading -n liqo-demo -o yaml
```

```
...
status:
  offloadingPhase: Ready
  remoteNamespaceName: liqo-demo
  remoteNamespacesConditions:
    florence-539f8b:
      - lastTransitionTime: "2022-05-05T15:08:33Z"
        message: You have not selected this cluster through ClusterSelector fields
        reason: ClusterNotSelected
```

(continues on next page)

(continued from previous page)

```

status: "False"
type: OffloadingRequired
naples-7881be:
- lastTransitionTime: "2022-05-05T15:08:43Z"
  message: The remote cluster has been selected through the ClusterSelector field
  reason: ClusterSelected
  status: "True"
  type: OffloadingRequired
- lastTransitionTime: "2022-05-05T15:08:43Z"
  message: Namespace correctly offloaded on this cluster
  reason: RemoteNamespaceCreated
  status: "True"
  type: Ready

```

Indeed, if you query for the namespaces in the *Naples* cluster, you should see the following output, confirming that the remote namespace has been correctly created by Liqo:

```
kubectl get namespaces liqo-demo --kubeconfig="$KUBECONFIG_NAPLES"
```

NAME	STATUS	AGE
liqo-demo	Active	70s

Instead, the same command executed in the *Florence* cluster should return an error, as the namespace has not been replicated:

```
kubectl get namespaces liqo-demo --kubeconfig="$KUBECONFIG_FLORENCE"
```

```
Error from server (NotFound): namespaces "liqo-demo" not found
```

13.4 Deploy applications

All constraints specified during namespace offloading are automatically enforced by Liqo, and merged with other pod-level specifications.

To verify this, you can now create two deployments in the `liqo-demo` namespace, characterized by additional *NodeAffinity* constraints. More precisely, one (`app-south`) is forced to be scheduled onto the virtual node representing the *Naples* cluster, while the other (`app-center`) is forced onto the *Florence* virtual cluster (which is incompatible with the namespace-level constraints).

```
kubectl apply -f ./manifests/deploy.yaml -n liqo-demo
```

Checking the pod status, it is possible to verify that one has been scheduled onto the *Naples* cluster, and it is correctly running, while the other remained *Pending* due to conflicting requirements (i.e., no node is available to satisfy all its constraints).

```
kubectl get pod -n liqo-demo -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
→ NOMINATED NODE	READINESS GATES					
app-center-6f6bd7547b-dh4m4	0/1	Pending	0	2m30s	<none>	<none>
→ <none>	<none>					

(continues on next page)

(continued from previous page)

app-south-85b9b99c4d-stwhw	1/1	Running	0	2m30s	10.204.0.12	liqo-
↔naples	<none>	<none>				

Note

You can remove the conflicting node affinity from the `app-center` deployment, and check that the generated pod gets scheduled onto either the *Venice* (i.e., locally) or the *Naples* cluster, as constrained by the namespace offloading configuration.

13.5 Tear down the playground

13.5.1 Unoffload namespaces

Before starting the uninstallation process, make sure that all namespaces are unoffloaded:

```
liqctl unoffload namespace liqo-demo
```

Every pod that was offloaded to a remote cluster is going to be rescheduled onto the local cluster.

13.5.2 Revoke peerings

Similarly, make sure that all the peerings are revoked:

```
liqctl unpeer out-of-band florence
liqctl unpeer out-of-band naples
```

At the end of the process, the virtual nodes are removed from the local cluster.

13.5.3 Uninstall Liqo

Now you can uninstall Liqo from your clusters with *liqctl*:

```
liqctl uninstall
liqctl uninstall --kubeconfig="$KUBECONFIG_FLORENCE"
liqctl uninstall --kubeconfig="$KUBECONFIG_NAPLES"
```

Purge

By default the Liqo CRDs will remain in the cluster, but they can be removed with the `--purge` flag:

```
liqctl uninstall --purge
liqctl uninstall --kubeconfig="$KUBECONFIG_FLORENCE" --purge
liqctl uninstall --kubeconfig="$KUBECONFIG_NAPLES" --purge
```

13.5.4 Destroy clusters

To teardown the KinD clusters, you can issue:

```
kind delete cluster --name venice  
kind delete cluster --name florence  
kind delete cluster --name naples
```


OFFLOADING A SERVICE

In this tutorial you will learn how to create a multi-cluster Service and how to consume it from each connected cluster. Specifically, you will deploy an application in a first cluster (*London*) and then offload the corresponding Service and transparently consume it from a second cluster (*New York*).

14.1 Provision the playground

First, check that you are compliant with the *requirements*.

Then, let's open a terminal on your machine and launch the following script, which creates the two above-mentioned clusters with KinD and installs Ligo on them. Each cluster is made by a single combined control-plane + worker node.

```
git clone https://github.com/liqotech/liqo.git
cd liqo
git checkout v0.6.0
cd examples/service-offloading
./setup.sh
```

Export the kubeconfigs environment variables to use them in the rest of the tutorial:

```
export KUBECONFIG="$PWD/liqo_kubeconf_london"
export KUBECONFIG_NEWYORK="$PWD/liqo_kubeconf_newyork"
```

Note

We suggest exporting the kubeconfig of the first cluster as default (i.e., KUBECONFIG), since it will be the entry point of the virtual cluster and you will mainly interact with it.

At this point, you should have two clusters with Ligo installed on them. The setup script named them **london** and **newyork**.

14.2 Peer the clusters

Once Liqo is installed in your clusters, you can establish new *peerings*. In this example, since the two API Servers are mutually reachable, you will use the *out-of-band peering approach*.

To implement the desired scenario, let's first retrieve the *peer command* from the *New York* cluster:

```
PEER_NEW_YORK=$(liqoctl generate peer-command --only-command --kubeconfig $KUBECONFIG_
↳NEWYORK)
```

Then, establish the peering from the *London* cluster:

```
echo "$PEER_NEW_YORK" | bash
```

When the above command returns successfully, you can check the peering status by running:

```
kubectl get foreignclusters
```

The output should look like the following, indicating that an outgoing peering is currently active towards the *New York* cluster, as well as that the cross-cluster network tunnel has been established:

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	AUTHENTICATION
↳ AGE					
newyork	OutOfBand	Established	None	Established	Established
↳ 61s					

14.3 Offload a service

Now, let's deploy a simple application composed of a *Deployment* and a *Service* in the *London* cluster.

First, you should create a hosting namespace in the *London* cluster:

```
kubectl create namespace liqo-demo
```

Then, deploy the application in the *London* cluster:

```
kubectl apply -f manifests/app.yaml -n liqo-demo
```

At this moment, you have an HTTP application serving JSON data through a *Service*, and running in the *London* cluster (i.e., locally). If you look at the *New York* cluster, you will not see the application yet.

To make it visible, you need to enable the Liqo offloading of the *Services* in the desired namespace to the *New York* cluster:

```
liqoctl offload namespace liqo-demo \
  --namespace-mapping-strategy EnforceSameName \
  --pod-offloading-strategy Local
```

This command enables the offloading of the *Services* in the *London* cluster to the *New York* cluster and sets:

- the namespace mapping strategy to *EnforceSameName*, which means that the namespace in the remote cluster is created with the same name as of the local one. This is particularly useful when you want to consume the *Services* in the remote cluster using the Kubernetes DNS service discovery (i.e. with `svc-name.namespace-name.svc.cluster.local`).

- the pod offloading strategy to *Local*, which means that the pods running in this namespace will be kept local and not scheduled on virtual nodes (i.e., no pod is offloaded to remote clusters).

Refer to the dedicated [usage page](#) for additional information concerning namespace offloading configurations.

Some seconds later, you should see that the *Service* has been replicated by the *resource reflection process*, and is now available in the *New York* cluster:

```
kubectl get services --namespace liqo-demo --kubeconfig "$KUBECONFIG_NEWYORK"
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE	SELECTOR
flights-service	ClusterIP	10.95.177.4	<none>	7999/TCP	14s	run=flights-
↪service						

The Service is characterized by a different *ClusterIP* address in the two clusters, since each cluster handles it independently. Additionally, you can also check that there is no application pod running in the *New York* cluster:

```
kubectl get pods --namespace liqo-demo --kubeconfig "$KUBECONFIG_NEWYORK"
```

```
No resources found in liqo-demo namespace.
```

14.3.1 Consume the service

Let's now consume the Service from both clusters from a different pod (e.g., a temporary shell).

Starting from the *London* cluster:

```
kubectl run consumer --rm -i --tty --image dwdraju/alpine-curl-jq -- /bin/sh
```

When the shell is ready, you can access the Service with `curl`:

```
curl -s -H 'accept: application/json' http://flights-service.liqo-demo:7999/schedule | jq .
```

A similar result is obtained executing the same command in a shell running in the *New York* cluster, although the backend pod is effectively running in the *London* cluster:

```
kubectl run consumer --rm -i --tty --image dwdraju/alpine-curl-jq \
  --kubeconfig $KUBECONFIG_NEWYORK -- /bin/sh
```

```
curl -s -H 'accept: application/json' http://flights-service.liqo-demo:7999/schedule | jq .
```

This quick example demonstrated how Liqo can **upgrade *ClusterIP* Services to multi-cluster Services**, allowing your local pods to transparently serve traffic originating from remote clusters with no additional configuration neither in the local cluster and/or applications nor in the remote ones.

14.4 Tear down the playground

14.4.1 Unoffload namespaces

Before starting the uninstallation process, make sure that all namespaces are unoffloaded:

```
liqctl unoffload namespace liqo-demo
```

Every pod that was offloaded to a remote cluster is going to be rescheduled onto the local cluster.

14.4.2 Revoke peerings

Similarly, make sure that all the peerings are revoked:

```
liqctl unpeer out-of-band newyork
```

At the end of the process, the virtual node is removed from the local cluster.

14.4.3 Uninstall Liqo

Now you can uninstall Liqo from your clusters with *liqctl*:

```
liqctl uninstall  
liqctl uninstall --kubeconfig="$KUBECONFIG_NEWYORK"
```

Purge

By default the Liqo CRDs will remain in the cluster, but they can be removed with the `--purge` flag:

```
liqctl uninstall --purge  
liqctl uninstall --kubeconfig="$KUBECONFIG_NEWYORK" --purge
```

14.4.4 Destroy clusters

To teardown the KinD clusters, you can issue:

```
kind delete cluster --name london  
kind delete cluster --name newyork
```

STATEFUL APPLICATIONS

This tutorial demonstrates how to use the core Liko features to deploy stateful applications. In particular, you will deploy a multi-master *mariadb-galera* database across a multi-cluster environment (composed of two clusters, respectively identified as *Turin* and *Lyon*), hence replicating the data in multiple regions.

15.1 Provision the playground

First, check that you are compliant with the *requirements*.

Then, let's open a terminal on your machine and launch the following script, which creates the two above-mentioned clusters with KinD and installs Liko on them. Each cluster is made by a single combined control-plane + worker node.

```
git clone https://github.com/liqotech/liqo.git
cd liqo
git checkout v0.6.0
cd examples/stateful-applications
./setup.sh
```

Export the kubeconfigs environment variables to use them in the rest of the tutorial:

```
export KUBECONFIG="$PWD/liqo_kubeconf_turin"
export KUBECONFIG_LYON="$PWD/liqo_kubeconf_lyon"
```

Note

The install script creates two clusters with no overlapping pod CIDRs. This is required by the *mariadb-galera* application to work correctly. Given it needs to know the real IP of the connected masters, it will not work correctly when natting is enabled.

Note

We suggest exporting the kubeconfig of the first cluster as default (i.e., KUBECONFIG), since it will be the entry point of the virtual cluster and you will mainly interact with it.

15.2 Peer the clusters

Once Liqo is installed in your clusters, you can establish new *peerings*. In this example, since the two API Servers are mutually reachable, you will use the *out-of-band peering approach*.

To implement the desired scenario, let's first retrieve the *peer command* from the *Lyon* cluster:

```
PEER_LYON=$(liqoctl generate peer-command --only-command --kubeconfig $KUBECONFIG_LYON)
```

Then, establish the peering from the *Turin* cluster:

```
echo "$PEER_LYON" | bash
```

When the above command returns successfully, you can check the peering status by running:

```
kubectl get foreignclusters
```

The output should look like the following, indicating that an outgoing peering is currently active towards the *Lyon* cluster,, as well as that the cross-cluster network tunnel has been established:

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	AUTHENTICATION	
↪AGE						└─
lyon	OutOfBand	Established	None	Established	Established	└─
↪1m28s						

15.3 Deploy a stateful application

In this step, you will deploy a *mariadb-galera* database using the [Bitnami helm chart](#).

First, you need to add the helm repository:

```
helm repo add bitnami https://charts.bitnami.com/bitnami
```

Then, create the namespace and offload it to remote clusters:

```
kubectl create namespace liqo-demo
liqoctl offload namespace liqo-demo --namespace-mapping-strategy EnforceSameName
```

This command will create a twin *liqo-demo* namespace in the *Lyon* cluster. Refer to the dedicated [usage page](#) for additional information concerning namespace offloading configurations.

Now, deploy the helm chart:

```
helm install db bitnami/mariadb-galera -n liqo-demo -f manifests/values.yaml
```

The release is configured to:

- have two replicas;
- spread the replicas across the cluster (i.e., a hard pod anti-affinity is set);
- use the *liqo virtual storage class*.

Check that these constraints are met by typing:

```
kubectl get pods -n liqo-demo -o wide
```

After a while (the startup process might require a few minutes), you should see two replicas of a *StatefulSet* spread over two different nodes (one local and one remote).

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
↪		NOMINATED NODE	READINESS GATES			
db-mariadb-galera-0	1/1	Running	0	3m13s	10.210.0.15	liqo-lyon
↪		<none>	<none>			
db-mariadb-galera-1	1/1	Running	0	2m6s	10.200.0.17	turin-control-
↪plane		<none>	<none>			

15.4 Consume the database

When the database is up and running, check that it is operating as expected executing a simple SQL client in your cluster:

```
kubectl run db-mariadb-galera-client --rm --tty -i \
  --restart='Never' --namespace default \
  --image docker.io/bitnami/mariadb-galera:10.6.7-debian-10-r56 \
  --command \
  -- mysql -h db-mariadb-galera.liqo-demo -uuser -ppassword my_database
```

And then create an example table and insert some data:

```
CREATE TABLE People (
  PersonID int,
  LastName varchar(255),
  FirstName varchar(255),
  Address varchar(255),
  City varchar(255)
);

INSERT INTO People
(PersonID, LastName, FirstName, Address, City)
VALUES
(1, 'Smith', 'John', '123 Main St', 'Anytown');
```

You are now able to query the database and grab the data:

```
SELECT * FROM People;
```

```
+-----+-----+-----+-----+-----+
| PersonID | LastName | FirstName | Address   | City     |
+-----+-----+-----+-----+-----+
|         1 | Smith   | John     | 123 Main St | Anytown |
+-----+-----+-----+-----+-----+
1 row in set (0.000 sec)
```

15.4.1 Database failures toleration

With this setup the applications running on a cluster can tolerate failures of the local database replica.

This can be checked deleting one of the replicas:

```
kubect1 delete pod db-mariadb-galera-0 -n liqo-demo
```

And querying again for your data:

```
kubect1 run db-mariadb-galera-client --rm --tty -i \  
  --restart='Never' --namespace default \  
  --image docker.io/bitnami/mariadb-galera:10.6.7-debian-10-r56 \  
  --command \  
  -- mysql -h db-mariadb-galera.liqo-demo -uuser -ppassword my_database \  
  --execute "SELECT * FROM People;"
```

Pro-tip

Try deleting the other replica and query again.

NOTE: at least one of the two replicas should be always running, be careful deleting all of them.

Note

You can run exactly the same commands to query the data from the other cluster, and you will get the same results.

15.5 Tear down the playground

15.5.1 Unoffload namespaces

Before starting the uninstallation process, make sure that all namespaces are unoffloaded:

```
liqoctl unoffload namespace liqo-demo
```

Every pod that was offloaded to a remote cluster is going to be rescheduled onto the local cluster.

15.5.2 Revoke peerings

Similarly, make sure that all the peerings are revoked:

```
liqoctl unpeer out-of-band lyon
```

At the end of the process, the virtual node is removed from the local cluster.

15.5.3 Uninstall Liqo

Now you can remove Liqo from your clusters with *liqoctl*:

```
liqoctl uninstall  
liqoctl uninstall --kubeconfig="$KUBECONFIG_LYON"
```

Purge

By default the Liqo CRDs will remain in the cluster, but they can be removed with the `--purge` flag:

```
liqoctl uninstall --purge  
liqoctl uninstall --kubeconfig="$KUBECONFIG_LYON" --purge
```

15.5.4 Destroy clusters

To teardown the KinD clusters, you can issue:

```
kind delete cluster --name turin  
kind delete cluster --name lyon
```


GLOBAL INGRESS

In this tutorial, you will learn how to leverage Liqo and K8GB to deploy and expose a multi-cluster application through a *global ingress*. More in detail, this enables improved load balancing and distribution of the external traffic towards the application replicated across multiple clusters.

The figure below outlines the high-level scenario, with a client consuming an application from either cluster 1 (e.g., located in EU) or cluster 2 (e.g., located in the US), based on the endpoint returned by the DNS server.

16.1 Provision the playground

First, check that you are compliant with the *requirements*. Additionally, this example requires k3d to be installed in your system. Specifically, this tool is leveraged instead of KinD to match the K8GB Sample Demo.

To provision the playground, clone the *Liqo repository* and run the setup script:

```
git clone https://github.com/liqotech/liqo.git
cd liqo
git checkout v0.6.0
cd examples/global-ingress
./setup.sh
```

The setup script creates three k3s clusters and deploys the appropriate infrastructural application on top of them, as detailed in the following:

- **edgedns**: this cluster will be used to deploy the DNS service. In a production environment, this should be an external DNS service (e.g. AWS Route53). It includes the Bind Server (manifests in `manifests/edge` folder).
- **gslb-eu** and **gslb-us**: these clusters will be used to deploy the application. They include:
 - **ExternalDNS**: it is responsible for configuring the DNS entries.
 - **Ingress Nginx**: it is responsible for handling the local ingress traffic.
 - **K8GB**: it configures the multi-cluster ingress.
 - **Liqo**: it enables the application to spread across multiple clusters, and takes care of reflecting the required resources.

Export the kubeconfigs environment variables to use them in the rest of the tutorial:

```
export KUBECONFIG_DNS=$(k3d kubeconfig write edgedns)
export KUBECONFIG=$(k3d kubeconfig write gslb-eu)
export KUBECONFIG_US=$(k3d kubeconfig write gslb-us)
```

Note

We suggest exporting the kubeconfig of the *gslb-eu* as default (i.e., KUBECONFIG), since it will be the entry point of the virtual cluster and you will mainly interact with it.

16.2 Peer the clusters

Once Liqo is installed in your clusters, you can establish new *peerings*. In this example, since the two API Servers are mutually reachable, you will use the *out-of-band peering approach*.

Specifically, to implement the desired scenario, you should enable a peering from the *gslb-eu* cluster to the *gslb-us* cluster. This will allow Liqo to *offload workloads and reflect services* from the first cluster to the second cluster.

To proceed, first generate a new *peer command* from the *gslb-us* cluster:

```
PEER_US=$(liqctl generate peer-command --only-command --kubeconfig $KUBECONFIG_US)
```

And then, run the generated command from the *gslb-eu* cluster:

```
echo "$PEER_US" | bash
```

When the above command returns successfully, you can check the peering status by running:

```
kubectl get foreignclusters
```

The output should look like the following, indicating that an outgoing peering is currently active towards the *gslb-us* cluster, as well as that the cross-cluster network tunnel has been established:

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	AUTHENTICATION
↪ AGE					
gslb-us	OutOfBand	Established	None	Established	Established
↪ 57s					

Additionally, you should see a new virtual node (*liqo-gslb-us*) in the *gslb-eu* cluster, and representing the whole *gslb-us* cluster. Every pod scheduled onto this node will be automatically offloaded to the remote cluster by Liqo.

```
kubectl get node --selector=liqo.io/type=virtual-node
```

The output should be similar to:

NAME	STATUS	ROLES	AGE	VERSION
liqo-gslb-us	Ready	agent	17s	v1.22.6+k3s1

16.3 Deploy an application

Now that the Liqo peering is established, and the virtual node is ready, it is possible to proceed deploying the *podinfo* demo application. This application serves a web-page showing different information, including the name of the pod; hence, easily identifying which replica is generating the HTTP response.

First, create a hosting namespace in the *gslb-eu* cluster, and offload it to the remote cluster through Liqo.

```
kubectl create namespace podinfo
liqctl offload namespace podinfo --namespace-mapping-strategy EnforceSameName
```

At this point, it is possible to deploy the *podinfo* helm chart in the *podinfo* namespace:

```
helm upgrade --install podinfo --namespace podinfo \
  podinfo/podinfo -f manifests/values/podinfo.yaml
```

This chart creates a *Deployment* with a *custom affinity* to ensure that the two frontend replicas are scheduled on different nodes and clusters:

```
affinity:
  nodeAffinity:
    requiredDuringSchedulingIgnoredDuringExecution:
      nodeSelectorTerms:
        - matchExpressions:
            - key: node-role.kubernetes.io/control-plane
              operator: DoesNotExist
  podAntiAffinity:
    requiredDuringSchedulingIgnoredDuringExecution:
      - labelSelector:
          matchExpressions:
            - key: app.kubernetes.io/name
              operator: In
            values:
              - podinfo
        topologyKey: "kubernetes.io/hostname"
```

Additionally, it creates an *Ingress* resource configured with the `k8gb.io/strategy: roundRobin` annotation. This annotation will instruct the [K8GB Global Ingress Controller](#) to distribute the traffic across the different clusters.

16.4 Check application spreading

Let's now check that Liqo replicated the ingress resource in both clusters and that each *Nginx Ingress Controller* was able to assign them the correct IPs (different for each cluster).

Note

You can see the output for the second cluster appending the `--kubeconfig $KUBECONFIG_US` flag to each command.

```
kubectl get ingress -n podinfo
```

The output in the *gslb-eu* cluster should be similar to:

NAME	CLASS	HOSTS	ADDRESS	PORTS	AGE
podinfo	nginx	liqo.cloud.example.com	172.19.0.3,172.19.0.4	80	6m9s

While the output in the *gslb-us* cluster should be similar to:

NAME	CLASS	HOSTS	ADDRESS	PORTS	AGE
podinfo	nginx	liqo.cloud.example.com	172.19.0.5,172.19.0.6	80	6m16s

With reference to the output above, the `liqo.cloud.example.com` hostname is served in the demo environment on:

- 172.19.0.3, 172.19.0.4: addresses exposed by cluster *gslb-eu*
- 172.19.0.5, 172.19.0.6: addresses exposed by cluster *gslb-us*

Each local *K8GB* installation creates a *Gslb* resource with the Ingress information and the given strategy (*RoundRobin* in this case), and *ExternalDNS* populates the DNS records accordingly.

On the *gslb-eu* cluster, the command:

```
kubectl get gslbs.k8gb.absa.oss -n podinfo podinfo -o yaml
```

should return an output along the lines of:

```
apiVersion: k8gb.absa.oss/v1beta1
kind: Gslb
metadata:
  annotations:
    k8gb.io/strategy: roundRobin
  name: podinfo
  namespace: podinfo
spec:
  ingress:
    ingressClassName: nginx
    rules:
    - host: liqo.cloud.example.com
      http:
        paths:
        - backend:
            service:
              name: podinfo
              port:
                number: 9898
          path: /
          pathType: ImplementationSpecific
  strategy:
    dnsTtlSeconds: 30
    splitBrainThresholdSeconds: 300
    type: roundRobin
status:
  geoTag: eu
  healthyRecords:
    liqo.cloud.example.com:
    - 172.19.0.3
    - 172.19.0.4
    - 172.19.0.5
```

(continues on next page)

(continued from previous page)

```
- 172.19.0.6
serviceHealth:
  liqo.cloud.example.com: Healthy
```

Similarly, when issuing the command from the *gslb-us* cluster:

```
kubectl get gslbs.k8gb.absa.oss -n podinfo podinfo -o yaml --kubeconfig $KUBECONFIG_US
```

```
apiVersion: k8gb.absa.oss/v1beta1
kind: Gslb
metadata:
  annotations:
    k8gb.io/strategy: roundRobin
  name: podinfo
  namespace: podinfo
spec:
  ingress:
    ingressClassName: nginx
    rules:
      - host: liqo.cloud.example.com
        http:
          paths:
            - backend:
                service:
                  name: podinfo
                  port:
                    number: 9898
              path: /
              pathType: ImplementationSpecific
    strategy:
      dnsTtlSeconds: 30
      splitBrainThresholdSeconds: 300
      type: roundRobin
status:
  geoTag: us
  healthyRecords:
    liqo.cloud.example.com:
      - 172.19.0.5
      - 172.19.0.6
      - 172.19.0.3
      - 172.19.0.4
  serviceHealth:
    liqo.cloud.example.com: Healthy
```

In both clusters, the *Gslb* resources are pretty identical; they only differ for the *geoTag* field. The resource status also reports:

- the *serviceHealth* status, that should be *Healthy* for both clusters
- the list of IPs exposing the HTTP service: they are the IPs of the nodes of both clusters since the *Nginx Ingress Controller* is deployed in *HostNetwork DaemonSet* mode.

16.5 Check service reachability

Since *podinfo* is an HTTP service, you can contact it using the *curl* command. Use the *-v* option to understand which of the nodes is being targeted.

You need to use the DNS server in order to resolve the hostname to the IP address of the service. To this end, create a pod in one of the clusters (it does not matter which one) overriding its DNS configuration.

```
HOSTNAME="liqo.cloud.example.com"
K8GB_COREDNS_IP=$(kubectl get svc k8gb-coredns -n k8gb -o custom-columns='IP:spec.
↪clusterIP' --no-headers)

kubectl run -it --rm curl --restart=Never --image=curlimages/curl:7.82.0 --command \
  --overrides "{\"spec\":{\"dnsConfig\":{\"nameservers\":[\"${K8GB_COREDNS_IP}\"]},\
↪\"dnsPolicy\":\"None\"}}" \
  -- curl $HOSTNAME -v
```

Note

Launching this pod several times, you will see different IPs and different frontend pods answering in a round-robin fashion (as set in the *Gslb* policy).

```
* Trying 172.19.0.3:80...
* Connected to liqo.cloud.example.com (172.19.0.3) port 80 (#0)
...
{
  "hostname": "podinfo-67f46d9b5f-xrbmg",
  "version": "6.1.4",
  "revision": "",
  ...
}
```

```
* Trying 172.19.0.6:80...
* Connected to liqo.cloud.example.com (172.19.0.6) port 80 (#0)
...
{
  "hostname": "podinfo-67f46d9b5f-xrbmg",
  "version": "6.1.4",
  "revision": "",
  ...
}
```

```
* Trying 172.19.0.3:80...
* Connected to liqo.cloud.example.com (172.19.0.3) port 80 (#0)
...
{
  "hostname": "podinfo-67f46d9b5f-cmnp5",
  "version": "6.1.4",
  "revision": "",
  ...
}
```

This brief tutorial showed how you could leverage Liqo and *K8GB* to deploy and expose a multi-cluster application. In addition to the *RoundRobin* policy, which provides load distribution among clusters, *K8GB* allows favoring closer

endpoints (through the *GeoIP* strategy), or adopt a *Failover* policy. Additional details are provided in its [official documentation](#).

16.6 Tear down the playground

16.6.1 Unoffload namespaces

Before starting the uninstallation process, make sure that all namespaces are unoffloaded:

```
liqctl unoffload namespace podinfo
```

Every pod that was offloaded to a remote cluster is going to be rescheduled onto the local cluster.

16.6.2 Revoke peerings

Similarly, make sure that all the peerings are revoked:

```
liqctl unpeer out-of-band gslb-us
```

At the end of the process, the virtual node is removed from the local cluster.

16.6.3 Uninstall Liqo

Now you can remove Liqo from your clusters with *liqctl*:

```
liqctl uninstall  
liqctl uninstall --kubeconfig="$KUBECONFIG_US"
```

Purge

By default the Liqo CRDs will remain in the cluster, but they can be removed with the `--purge` flag:

```
liqctl uninstall --purge  
liqctl uninstall --kubeconfig="$KUBECONFIG_US" --purge
```

16.6.4 Destroy clusters

To teardown the k3d clusters, you can issue:

```
k3d cluster delete gslb-eu gslb-us edgedns
```


PEER TWO CLUSTERS

This section describes the procedure to **establish a peering** with a remote cluster, using one of the two alternative approaches featured by Ligo. You can refer to the *dedicated features section* for a high-level presentation of their characteristics, and the associated trade-offs.

Warning: The establishment of a peering with a remote cluster leveraging a **different version of Ligo**, net of patch releases, is currently **not supported**, and could lead to unexpected results.

17.1 Overview

The peering process leverages *liqoctl* to interact with the clusters, abstracting the creation and update of the appropriate custom resources. To this end, the most important one is the *ForeignCluster* resource, which **represents a remote cluster**, including its identity, the associated authentication endpoint, and the desired peering state (i.e., whether it should be established, and in which directions). Additionally, its status reports a **summary of the current peering status**, detailing whether the different phases (e.g., authentication, network establishment, resource negotiation, ...) correctly succeeded.

The following sections present the respective procedures to **peer a local cluster A** (i.e., the *consumer*), with a **remote cluster B** (i.e., the *provider*). At the end of the process, a new **virtual node** is created in the consumer, abstracting the resources shared by the provider, and enabling seamless **pod offloading** to the remote cluster. Additional details are also provided to enable the reverse peering direction, hence achieving a **bidirectional peering**, allowing both clusters to offload a part of their workloads to the other.

All examples leverage two different *contexts* to refer to *consumer* and *provider* clusters, respectively named *consumer* and *provider*.

Note

liqoctl displays a *kubectl* compatible behavior concerning Kubernetes API access, hence supporting the KUBECONFIG environment variable, as well as all the standard flags, including `--kubeconfig` and `--context`. Ensure you selected the correct target cluster before issuing *liqoctl* commands (as you would do with *kubectl*).

17.2 Out-of-band control plane

Briefly, the procedure to establish an *out-of-band control plane peering* consists of a first step performed on the *provider*, to **retrieve the set of information** required (i.e., authentication endpoint and token, cluster ID, ...), followed by the creation, on the *consumer*, of the necessary resources to **start the actual peering**. The remainder of the process, including identity retrieval, resource negotiation and network tunnel establishment is **performed automatically** by Liqo, through a mutual exchange of information and negotiation between the two clusters involved.

17.2.1 Information retrieval

To proceed, ensure that you are operating in the *provider* cluster, and then issue the `liqoctl generate peer-command` command:

```
liqoctl --context=provider generate peer-command
```

This retrieves the information concerning the *provider* cluster (i.e., authentication endpoint and token, cluster ID, ...) and generates a command that can be executed on a *different* cluster (i.e., the *consumer*) to establish an out-of-band outgoing peering towards the *provider* cluster.

An example of the resulting command is the following:

```
liqoctl peer out-of-band <cluster-name> --auth-url <auth-url> \  
--cluster-id <cluster-id> --auth-token <auth-token>
```

17.2.2 Peering establishment

Once obtained the peering command, it is possible to execute it in the *consumer* cluster, to kick off the peering process.

Warning: Pay attention to operate in the correct cluster, possibly adding the appropriate flags to the generated command (e.g., `--context=consumer`).

```
liqoctl --context=consumer peer out-of-band <cluster-name> --auth-url <auth-url> \  
--cluster-id <cluster-id> --auth-token <auth-token>
```

The above command configures the appropriate authentication token, and then creates a new *ForeignCluster* resource in the *consumer cluster*. Finally, it waits for the different peering phases to complete (this might require a few seconds, depending on the download time of the Liqo virtual kubelet image).

The *ForeignCluster* resource can be inspected through `kubect`:

```
kubectl --context=consumer get foreignclusters
```

If the peering process completed successfully, you should observe an output similar to the following, indicating that the cross-cluster network tunnel has been established, and an outgoing peering is currently active (i.e., the *consumer* cluster can offload workloads to the *provider* one, but not vice versa):

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	AUTHENTICATION
provider	OutOfBand	Established	None	Established	Established

At the same time, a new *virtual node* should have been created in the *consumer* cluster. Specifically:

```
kubectl --context=consumer get nodes -l liqo.io/type=virtual-node
```

Should return an output similar to the following:

NAME	STATUS	ROLES	AGE	VERSION
liqo-provider	Ready	agent	179m	v1.23.4

Note

The name of the *ForeignCluster* resource, as well as that of the *virtual node*, reflects the cluster name specified with the *liqoctl peer out-of-band* command.

17.2.3 Bidirectional peering

Once the peering from the *consumer* to the *provider* has been established, the reverse direction (i.e., leading to a bidirectional peering) can be enabled through a simpler command, since the *ForeignCluster* resource is already present:

```
liqoctl --context=provider peer consumer
```

17.2.4 Tear down

An out-of-band peering can be disabled leveraging the symmetric *liqoctl unpeer* command, causing the local virtual node (abstracting the remote cluster) to be destroyed, and all offloaded workloads to be rescheduled:

```
liqoctl --context=consumer unpeer out-of-band
```

Note

The reverse peering direction, if any, is preserved, and the remote cluster can continue offloading workloads to its virtual node representing the local cluster. In this case, the command *emits a warning*, and it does not proceed deleting the *ForeignCluster* resource. Hence, the same command shall be executed on both clusters to completely tear down a bidirectional peering.

In case only one peering direction shall be teared down, while preserving the opposite, it is suggested to leverage the appropriate *liqoctl unpeer* command to disable the outgoing peering (e.g., on the *provider* cluster):

```
liqoctl --context=provider unpeer consumer
```

17.3 In-band control plane

Briefly, the procedure to establish an *in-band control plane peering* consists of a first step performed by *liqoctl*, which interacts alternatively with both clusters to **establish the cross-cluster VPN tunnel**, exchange the **authentication tokens** and configure the Liqo control plane traffic to flow inside the VPN. The remainder of the process, including identity retrieval and resource negotiation, is **performed automatically** by Liqo, through a mutual exchange of information and negotiation between the two clusters involved.

Note

The host used to issue the `liqoctl peer in-band` command must have **concurrent access to both clusters** (i.e., *consumer* and *provider*) while carrying out the in-band control plane peering process. To this end, these subcommands feature a parallel set of flags concerning Kubernetes API access to the remote cluster, in the form `--remote-<flag>` (e.g., `--remote-kubeconfig`, `--remote-context`).

17.3.1 Peering establishment

The in-band control plane peering process can be started leveraging a single `liqoctl` command:

```
liqoctl peer in-band --context=consumer --remote-context=provider
```

The above command outputs a set of information concerning the different operations performed on the two clusters. Notably, it exchanges the appropriate authentication tokens, establishes the cross-cluster VPN tunnel, and then creates a new *ForeignCluster* resource in *both clusters*. Finally, it waits for the different peering phases to complete (this might require a few seconds, depending on the download time of the Liqo virtual kubelet image).

The *ForeignCluster* resource can be inspected through `kubectl` (e.g., on the *consumer*):

```
kubectl --context=consumer get foreignclusters
```

If the peering process completed successfully, you should observe an output similar to the following, indicating that the cross-cluster network tunnel has been established, and an outgoing peering is currently active (i.e., the *consumer* cluster can offload workloads to the *provider* one, but not vice versa):

NAME	TYPE	OUTGOING PEERING	INCOMING PEERING	NETWORKING	AUTHENTICATION
provider	InBand	Established	None	Established	Established

At the same time, a new *virtual node* should have been created in the *consumer* cluster. Specifically:

```
kubectl --context=consumer get nodes -l liqo.io/type=virtual-node
```

Should return an output similar to the following:

NAME	STATUS	ROLES	AGE	VERSION
liqo-provider	Ready	agent	179m	v1.23.4

Note

The name of the *ForeignCluster* resource, as well as that of the *virtual node*, reflects the cluster name specified by the remote cluster administrators at install time.

17.3.2 Bidirectional peering

A bidirectional in-band peering can be established adding the `--bidirectional` flag to the `liqoctl peer` command invocation:

```
liqoctl peer in-band --context=consumer --remote-context=provider --bidirectional
```

Note

The `liqctl peer in-band` command is idempotent, and can be re-executed without side effects to enable a bidirectional peering.

Alternatively, the reverse peering can be also activated executing the following on the *provider* cluster:

```
liqctl --context=provider peer consumer
```

17.3.3 Tear down

An in-band peering can be disabled leveraging the symmetric `liqctl unpeer` command, causing both virtual nodes (if present) to be destroyed, all offloaded workloads to be rescheduled, and finally tearing down the cross-cluster VPN tunnel:

```
liqctl unpeer in-band --context=consumer --remote-context=provider
```

In case only one peering direction shall be teared down, while preserving the opposite, it is possible to leverage the appropriate `liqctl unpeer` command to disable the outgoing peering (e.g., on the *provider* cluster):

```
liqctl --context=provider unpeer consumer
```


NAMESPACE OFFLOADING

This section presents the operational procedure to **offload a namespace** to (possibly a subset of) the **remote clusters** peered with the local cluster. Hence, enabling **pod offloading**, as well as triggering the *resource reflection* process: additional details about namespace extension in Ligo are provided in the dedicated *namespace extension features section*.

18.1 Overview

The offloading of a namespace can be easily controlled through the dedicated *liqctl* commands, which abstract the creation and update of the appropriate custom resources. In this context, the most important one is the *NamespaceOffloading* resource, which enables the offloading of the corresponding namespace, configuring at the same time the subset of target remote clusters, additional constraints concerning pod offloading and the naming strategy. Moreover, different namespaces can be characterized by different configurations, hence achieving a high degree of flexibility. Finally, the *NamespaceOffloading* status reports for each remote cluster a **summary about its status** (i.e., whether the remote cluster has been selected for offloading, and the twin namespace has been correctly created).

18.2 Offloading a namespace

A given namespace *foo* can be offloaded, leveraging the default configuration, through:

```
liqctl offload namespace foo
```

Alternatively, the underlying *NamespaceOffloading* resource can be generated and output (either in *yaml* or *json* format) leveraging the dedicated `--output` flag:

```
liqctl offload namespace foo --output yaml
```

Then, the resulting manifest can be applied with *kubectl*, or through automation tools (e.g., by means of GitOps approaches).

Note

Possible race conditions might occur in case a *NamespaceOffloading* resource is created at the same time (e.g., as a batch) as pods (or higher level abstractions such as *Deployments*), preventing them from being considered for offloading until the *NamespaceOffloading* resource is not processed.

This situation can be prevented manually labeling in advance the hosting namespace with the *liq.io/scheduling-enabled=true* label, hence enabling the Ligo mutating webhook and causing pod creations to be rejected until pod

offloading is possible. Still, this causes no problems, as the Kubernetes abstractions (e.g., *Deployments*) ensure that the desired pods get eventually created correctly.

Regardless of the approach adopted, namespace offloading can be further configured in terms of the three main parameters presented below, each one exposed through a dedicated CLI flag.

18.2.1 Namespace mapping strategy

The *namespace mapping strategy* defines the naming strategy used to create the remote namespaces, and can be configured through the `--namespace-mapping-strategy` flag. The accepted values are:

- **DefaultName** (default): to **prevent conflicts** on the target cluster, remote namespace names are generated as the concatenation of the local namespace name, the cluster name of the local cluster and a unique identifier (e.g., *foo* could be mapped to *foo-lively-voice-dd8531*).
- **EnforceSameName**: remote namespaces are named after the local cluster's namespace. This approach ensures **naming transparency**, which is required by certain applications, as well as guarantees that **cross-namespace DNS queries** referring to reflected services work out of the box (i.e., without adapting the target namespace name). Yet, it can lead to **conflicts** in case a namespace with the same name already exists inside the selected remote clusters, ultimately causing the remote namespace creation request to be rejected.

Note

Once configured for a given namespace, the *namespace mapping strategy* is **immutable**, and any modification is prevented by a dedicated Liqo webhook. In case a different strategy is desired, it is necessary to first *unoffload* the namespace, and then re-offload it with the new parameters.

18.2.2 Pod offloading strategy

The *pod offloading strategy* defines high-level constraints about pod scheduling, and can be configured through the `--pod-offloading-strategy` flag. The accepted values are:

- **LocalAndRemote** (default): pods deployed in the local namespace can be scheduled **both onto local nodes and onto virtual nodes**, hence possibly offloaded to remote clusters.
- **Local**: pods deployed in the local namespace are enforced to be scheduled onto **local nodes only**, hence never offloaded to remote clusters. The extension of a namespace, forcing at the same time all pods to be scheduled locally, enables the consumption of local services from the remote cluster, as shown in the [service offloading example](#).
- **Remote**: pods deployed in the local namespace are enforced to be scheduled onto **remote nodes only**, hence always offloaded to remote clusters.

Note

The *pod offloading strategy* applies to pods only, while the other objects that live in namespaces selected for offloading, and managed by the resource reflection process, are always replicated to (possibly a subset of) the remote clusters, as specified through the *cluster selector* (more details below).

Warning: Due to current limitations of Liqo, the pods violating the *pod offloading strategy* are not automatically evicted following an update of this policy to a more restrictive value (e.g., *LocalAndRemote* to *Remote*) after the initial creation.

18.2.3 Cluster selector

The *cluster selector* provides the possibility to **restrict the set of remote clusters** (in case more than one peering is active) selected as targets for offloading the given namespace. The *twin* namespace is not created in clusters that do not match the cluster selector, as well as the resource reflection mechanism is not activated for those namespaces. Yet, different *cluster selectors* can be specified for different namespaces, depending on the desired configuration.

The cluster selector follows the standard **label selector** syntax, and refers to the Kubernetes labels characterizing the **virtual nodes**. Specifically, these include both the set of labels suggested by the remote cluster during the peering process and automatically propagated by Liqo, as well as possible additional ones added by the local cluster administrators.

The cluster selector can be expressed through the `--selector` flag, which can be optionally repeated multiple times to specify alternative requirements (i.e., in logical OR). For instance:

- `--selector 'region in (europe,us-west), !staging'` would match all clusters located in the *europe* or *us-west* region, *AND* not including the *staging* label.
- `--selector 'region in (europe,us-west)' --selector '!staging'` would match all clusters located in the *europe* or *us-west* region, *OR* not including the *staging* label.

In case no *cluster selector* is specified, all remote clusters are selected as targets for namespace offloading. In other words, an empty *cluster selector* matches all virtual clusters.

18.3 Unoffloading a namespace

The offloading of a namespace can be disabled through the dedicated *liqoctl* command, causing in turn the deletion of all resources reflected to remote clusters (including the namespaces themselves), and triggering the rescheduling of all offloaded pods locally:

```
liqoctl unoffload namespace foo
```

Warning: Disabling the offloading of a namespace is a **destructive operation**, since all resources created in remote namespaces (either automatically or manually) get removed, including possible **persistent storage volumes**. Before proceeding, double-check that the correct namespace has been selected, and ensure no important data is still present.

RESOURCE REFLECTION

This section characterizes the *resource reflection* process (including also *pod offloading*), detailing how the different resources are propagated to remote clusters and which fields are mutated.

Briefly, the set of supported resources includes (by category):

- **Workload:** *Pods*
- **Exposition:** *Services, EndpointSlices, Ingresses*
- **Storage:** *PersistentVolumeClaims, PersistentVolumes*
- **Configuration:** *ConfigMaps, Secrets*

Note

The reflection of a given object belonging to the *Exposition* or *Configuration* categories, and living in a namespace enabled for offloading, can be manually disabled adding the `liqo.io/skip-reflection` annotation to the object itself.

Additionally, the reflection of a given type of resources (e.g., *Secrets*) towards remote clusters can be completely disabled setting the corresponding `--<resource>-reflection-workers=0` virtual kubelet flag at install time:

```
liqoctl install ... --set "virtualKubelet.extra.args={--secret-reflection-workers=0}"
```

19.1 Pods offloading

Liqo leverages a custom resource, named *ShadowPod*, combined with an appropriate enforcement logic to ensure **remote pod resiliency** even in case of temporary connectivity loss between the local and remote clusters.

Pod specifications are propagated to the remote cluster **verbatim**, except for the following fields that are mutated:

- Removal of **scheduling constraints** (e.g., *Affinity, NodeSelector, SchedulerName, Preemption, ...*), as referring to the local cluster.
- Mutation of **service account** related information, to allow offloaded pods to transparently interact with the local (i.e., origin) API server, instead of the remote one.
- Enforcement of the properties concerning the usage of **host namespaces** (e.g., network, IPC, PID) to *false* (i.e., disabled), as potentially invasive and troublesome.

Note

Anti-affinity presets can be leveraged to specify predefined scheduling constraints for offloaded pods, spreading them across different nodes in the remote cluster. This feature is enabled through the `liqo.io/anti-affinity-preset` pod annotation, which can take three values:

- **propagate**: the anti-affinity constraints of the pod are propagated *verbatim* when offloaded to the remote cluster. Make sure that they match both the virtual node in the local cluster and at least one physical node in the remote cluster, otherwise the pod will fail to be scheduled (i.e., remain in pending status).
- **soft**: the pods sharing the same labels are *preferred* to be scheduled on different nodes (i.e., it is translated into a *preferredDuringSchedulingIgnoredDuringExecution* anti-affinity constraint).
- **hard**: the pods sharing the same labels are *required* to be scheduled on different nodes (i.e., it is translated into a *requiredDuringSchedulingIgnoredDuringExecution* anti-affinity constraint).

When set to *soft* or *hard*, the `liqo.io/anti-affinity-labels` annotation allows to select a subset of the pod label keys to build the anti-affinity constraints:

```
annotations:  
  liqo.io/anti-affinity-preset: soft  
  liqo.io/anti-affinity-labels: app.kubernetes.io/name,app.kubernetes.io/instance
```

Given that affinity constraints are *immutable*, the addition/removal of the annotations to/from an already existing pod *does not have any effect*. Make sure that the annotations are configured appropriately in the template of the managing object (e.g., *Deployment*, or *StatefulSet*).

Differently, **pod status** is propagated from the remote cluster to the local one, performing the following modifications:

- The *PodIP* is **remapped** according to the network fabric configuration, such as to be reachable from the other pods running in the same cluster.
- The *NodeIP* is replaced with the one of the corresponding virtual kubelet pod.
- The number of **container restarts** is augmented to account for the possible deletions of the remote pod (whose presence is enforced by the controlling *ShadowPod* resource).

Note

A pod living in a namespace not enabled for offloading, but manually forced to be scheduled in a virtual node, remains in *Pending* status, and it is signaled with the *OffloadingBackOff* reason. For instance, this can happen for system *DaemonSets* (e.g., CNI plugins), which tolerate all *taints* (hence, including the one associated with virtual nodes) and thus get scheduled on *all nodes*.

To prevent this behavior, it is necessary to explicitly modify the involved *DaemonSets*, adding a suitable *affinity* constraint excluding virtual nodes:

```
affinity:  
  nodeAffinity:  
    requiredDuringSchedulingIgnoredDuringExecution:  
      nodeSelectorTerms:  
        - matchExpressions:  
          - key: liqo.io/type  
            operator: NotIn  
            values:  
              - virtual-node
```

19.2 Service exposition

The reflection of **Service** and **EndpointSlice** resources is a key element to allow the seamless **intercommunication** between microservices spread across multiple clusters, enabling the usage of standard DNS discovery mechanisms. In addition, the propagation of **Ingresses** enables the definition of multiple points of entrance for the external traffic, especially when combined with additional tools such as **K8GB** (see the *global ingress example* for additional details).

19.2.1 Services

Services are reflected **verbatim** into remote clusters, except for what concerns the *ClusterIP*, *LoadBalancerIP* and *NodePort* fields (when applicable), which are left empty (hence defaulted by the remote cluster), as likely conflicting. Still, the usage of **standard DNS discovery** mechanisms (i.e., based on service name/namespace) abstracts away the *ClusterIP* differences, with each pod retrieving the correct IP address.

Note

In case *node port* correspondence across clusters is required, its propagation can be enforced adding the `liqo.io/force-remote-node-port=true` annotation to the involved service.

19.2.2 EndpointSlices

In the local cluster, **Services** are transparently handled by the vanilla Kubernetes control plane, since it has **full visibility of all pods** (even those offloaded), hence leading to the creation of the corresponding **EndpointSlice** entries. Differently, the control plane of each remote cluster perceives **only the pods running in that cluster**, and the standard *EndpointSlice* creation logic alone is not sufficient (as it would not include the pods hosted by other clusters).

This gap is filled by the Liqo **EndpointSlice reflection** logic, which takes care of propagating all *EndpointSlice* entries (i.e. endpoints) not already present in the destination cluster. During the propagation process, endpoint addresses are appropriately **remapped** according to the **network fabric** configuration, ensuring that the resulting IPs are reachable from the destination cluster.

Thanks to this approach, **multiple replicas** of the same microservice spread across different clusters, and backed by the same service, are handled transparently. Each pod, no matter where it is located, contributes with a distinct *EndpointSlice* entry, either by the standard control plane or through resource reflection, hence becoming eligible during the **Service load-balancing process**.

19.2.3 Ingresses

The propagation of **Ingress** resources enables the configuration of multiple points of entrance for **external traffic**. *Ingress* resources are propagated **verbatim** into remote clusters, except for the *IngressClassName* field, which is left empty. Hence, selecting the default *ingress class* in the remote cluster, as the local one (i.e., the one in the origin cluster) might not be present.

19.3 Persistent storage

The reflection of **PersistentVolumeClaims (PVCs)** and **PersistentVolumes (PVs)** is a key to enable the cross-cluster *Liqo storage fabric*. Specifically, the process is triggered when a PVC requiring the *Liqo storage class* is bound for the first time, and the requesting pod is scheduled in a virtual node (i.e., remote cluster). Upon this event, the **PVC is propagated verbatim** to the remote cluster, replacing the requested *StorageClass* with the one negotiated during the peering process.

Once created, the **resulting PV is reflected backwards** (i.e., from the remote to the local cluster), and the proper **affinity selectors** are added to **bind it to the virtual node**. Hence, subsequent pods mounting that *PV* will be scheduled on that virtual node, and eventually offloaded to the same remote cluster.

19.4 Configuration data

ConfigMaps and **Secrets** typically hold **configuration data** consumed by pods, and both types of resources are propagated by Liqo **verbatim** into remote clusters. In this respect, Liqo features also the propagation of **Secrets** holding **ServiceAccount tokens**, to enable offloaded pods to contact the Kubernetes API server of the origin cluster, as well as to support those applications leveraging *ServiceAccounts* for internal authentication purposes.

Warning: Currently, Liqo supports only the propagation of *ServiceAccount* tokens contained in the respective *Secret* object (i.e., *first party tokens*), and not of those to be retrieved from the *TokenRequest* API (i.e., *third party tokens*). Due to this limitation, service account reflection is currently *disabled* by default in Kubernetes v1.24+, as *ServiceAccounts* do not longer automatically generate the corresponding *Secret*.

STATEFUL APPLICATIONS

As introduced in the *storage fabric features section*, Ligo supports **multi-cluster stateful applications** by extending the classical approaches adopted in standard Kubernetes clusters.

20.1 Ligo virtual storage class

The Ligo virtual storage class is a *Storage Class* that embeds the logic to create the appropriate *PersistentVolumes*, depending on the target cluster the mounting pod is scheduled onto. All operations performed on virtual objects (i.e., *PersistentVolumeClaims (PVCs)* and *PersistentVolumes (PVs)* associated with the *liqo* storage class) are then automatically propagated by Ligo to the corresponding real ones (i.e., associated with the storage class available in the target cluster).

Additionally, once a real *PV* gets created, the corresponding virtual one is enriched with a set of policies to attract mounting pods in the appropriate cluster, following the **data gravity** approach.

The figure below shows an application pod consuming a virtual *PVC*, which in turn led to the creation of the associated virtual *PV*. This process is **completely transparent** from the management point of view, with the only difference being the name of the storage class.

<p>Warning: The deletion of the virtual <i>PVC</i> will cause the deletion of the real <i>PVC/PV</i>, and the stored data will be permanently lost.</p>

The Ligo control plane handles the binding of virtual *PVC* resources (i.e., associated with the *liqo* storage class) differently depending on the cluster where the mounting pod gets eventually scheduled onto, as detailed in the following.

20.1.1 Local cluster binding

In case a virtual *PVC* is bound to a pod initially **scheduled onto the local cluster** (i.e., a physical node), the Ligo control plane takes care of creating a twin *PVC* (in turn originating the corresponding twin *PV*) in the *liqo-storage* namespace, while mutating the *storage class* to that configured at Ligo installation time (with a fallback to the default one). A virtual *PV* is eventually created by Ligo to mirror the real one, effectively allowing pods to mount it and enforcing the *data gravity* constraints.

The resulting configuration is depicted in the figure below.

Current Limitations

Currently, the virtual storage class does not support the configuration of [Kubernetes mount options](#) and parameters.

20.1.2 Remote cluster binding

In case a virtual *PVC* is bound to a pod initially **scheduled onto a remote cluster** (i.e., a virtual node), the Liqo control plane takes care of creating a twin *PVC* (in turn originating the corresponding twin *PV*) in the *offloaded* namespace, while mutating the *storage class* to that negotiated at peering time (i.e., configured at Liqo installation time in the remote cluster, with a fallback to the default one). A virtual *PV* is eventually created by Liqo to mirror the real one, effectively allowing pods to mount it and enforcing the *data gravity* constraints.

The resulting configuration is depicted in the figure below.

Warning: The tearing down of the peering and/or the deletion of the offloaded namespace will cause the deletion of the real *PVC*, and the stored data will be **permanently lost**.

20.1.3 Move PVCs across clusters

Once a *PVC* is created in a given cluster, subsequent pods mounting that volume will be forced to be **scheduled onto the same cluster** to achieve storage locality, following the *data gravity* approach.

Still, if necessary, you can **manually move** the storage backing a virtual *PVC* (i.e., associated with the *liqo* storage class) from a cluster to another, leveraging the appropriate *liqoctl* command. Then, subsequent pods will get scheduled in the cluster the storage has been moved to.

Warning: This procedure requires the *PVC/PV* not to be bound to any pods during the entire process. In other words, live migration is currently not supported.

A given *PVC* can be moved to a target node (either physical, i.e., local, or virtual, i.e., remote) through the following command:

```
liqoctl move volume $PVC_NAME --namespace $NAMESPACE_NAME --node $TARGET_NODE_NAME
```

Where:

- `$PVC_NAME` is the name of the *PVC* to be moved.
- `$NAMESPACE_NAME` is the name of the namespace where the *PVC* lives in.
- `$TARGET_NODE_NAME` is the name of the node where the *PVC* will be moved to.

Under the hood, the migration process leverages the Liqo cross-cluster network fabric and the [Restic project](#) to back up the original data in a temporary repository, and then restore it in a brand-new *PVC* forced to be created in the target cluster.

Warning: *Liqo* and *liqoctl* **are not** backup tools. Make sure to properly back up important data before starting the migration process.

20.2 Externally managed storage

In addition to the virtual storage class, Liqo supports the offloading of pods that bind to ***cross-cluster storage managed by external solutions*** (e.g., managed by the cloud provider, or manually provisioned). Specifically, the *volumes* stanza of the pod specification is propagated verbatim to the offloaded pods, hence allowing to specify volumes available only remotely.

Note

In case a piece of externally managed storage is available only in one remote cluster, it is likely necessary to manually force pods to get scheduled exactly in that cluster. To prevent scheduling issues (e.g., the pod is marked as *Pending* since the local cluster has no visibility on the remote *PVC*), it is suggested to configure the target *nodeName* in the pod specifications to match that of the corresponding virtual nodes, hence bypassing the standard Kubernetes scheduling logic.

Warning: Due to current Liqo limitations, the remote namespace, including any *PVC* therein contained, will be **deleted** in case the local namespace is unoffloaded/deleted, or the peering is torn down.

PROMETHEUS METRICS

This section presents the metrics exposed by Ligo components, using the [Prometheus](#) format.

21.1 Scraping metrics

How metrics are scraped will depend on how you are operating your **Prometheus server(s)**. Ligo presumes that the [Prometheus Operator](#) is being used to run Prometheus, providing a **ServiceMonitor** resource for each component.

Metrics are **disabled** by default. In order to allow Prometheus to scrape metrics from the Ligo components, you should set the `--enable-metrics liqoctl` flag during installation (cf. [installation customization options](#)). This flag enables the metrics exposition and the **ServiceMonitor** resources creation.

Alternatively, it is possible to customize these settings using **Helm** (refer to the [Install with Helm](#) section for further details).


21.2 Cross-cluster network metrics

These metrics are available for each peered remote cluster, providing statistics about the cross-cluster network interconnections:

- **liqo_peer_receive_bytes_total**: the total number of bytes received from a remote cluster.
- **liqo_peer_transmit_bytes_total**: the total number of bytes transmitted to a remote cluster.
- **liqo_peer_latency_us**: the latency between the local cluster and a remote cluster.
- **liqo_peer_is_connected**: whether the network interconnection is established and works properly.

21.2.1 Grafana dashboard

We provide a [sample Grafana dashboard](#) to monitor the network interconnection of an arbitrary number of Ligo peerings. As presented in the screenshot below, it includes an overview section presenting the overall cross-cluster throughput, followed by detailed per-peering throughput and latency information.



CONTRIBUTING TO LIQO

First off, thank you for taking the time to contribute to Ligo!

This page lists a set of contributing guidelines, including suggestions about the local development of Ligo components and the execution of the automatic tests.

22.1 Repository structure

The Ligo repository structure follows the [Standard Go Project Layout](#).

22.2 Release notes generation

Ligo leverages the automatic release notes generation capabilities featured by GitHub. Specifically, PRs characterized by the following labels get included in the respective category:

- *kind/breaking*: Breaking Change
- *kind/feature*: New Features
- *kind/bug*: Bug Fixes
- *kind/cleanup*: Code Refactoring
- *kind/docs*: Documentation

22.3 Local development

While developing a new feature, it is typically useful to test the changes in a local environment, as well as debug the code to identify possible problems. To this end, you can leverage the *setup.sh* script provided for the *quick start example* to spawn two development clusters using **KinD**, and then install Ligo on both of them (you can refer to the *dedicated section* for additional information concerning the installation of development versions through `liqctl`):

```
./examples/quick-start/setup.sh
liqctl install kind --kubeconfig=./liqo_kubeconf_rome --version ...
liqctl install kind --kubeconfig=./liqo_kubeconf_milan --version ...
```

Once the environment is properly setup, it is possible to proceed according to one of the following approaches:

- Building and pushing the Docker image of the component under development to a registry, and appropriately editing the corresponding *Deployment/DaemonSet* to use the custom version. This allows to observe the modified

component in realistic conditions, and it is mandatory for the networking substratum, since it needs to interact with the underlying host configuration.

- Scaling to 0 the number of replicas of the component under development, copying its current configuration (i.e., command-line flags), and executing it locally (while targeting the appropriate cluster). This allows for faster development cycles, as well as for the usage of standard debugging techniques to troubleshoot possible issues.

22.4 Automatic tests

Liqo features two major test suites:

- *End-to-end (E2E) tests*, which assess the correct functioning of the main Liqo features.
- *Unit Tests*, which focus on each specific component, in multiple operating conditions.

Both test suites are automatically executed through the GitHub Actions pipelines, following the corresponding slash command trigger. A successful outcome is required to make PRs eligible for being considered for review and merged.

The following sections provide additional details concerning how to run the above tests in a local environment, for troubleshooting.

22.4.1 End-to-end tests

We suggest executing the E2E tests on a system with at least 8 GB of free RAM. Additionally, please review the requirements presented in the *Liqo examples section*, which also apply in this case (including the suggestions concerning increasing the maximum number of *inotify* watches).

Once all requirements are met, it is necessary to export the set of environment variables shown below, to configure the tests. In most scenarios, the only variable that needs to be modified is `LIQO_VERSION`, which should point to the SHA of the commit referring to the Liqo development version to be tested (the appropriate Docker images shall have been built in advance through the appropriate GitHub Actions pipeline).

```
export CLUSTER_NUMBER=4
export K8S_VERSION=v1.21.1
export CNI=kindnet
export TMPDIR=$(mktemp -d)
export BINDIR=${TMPDIR}/bin
export TEMPLATE_DIR=${PWD}/test/e2e/pipeline/infra/kind
export NAMESPACE=liqo
export KUBECONFIGDIR=${TMPDIR}/kubeconfigs
export LIQO_VERSION=<YOUR_COMMIT_ID>
export INFRA=kind
export LIQOCTL=${BINDIR}/liqoctl
export POD_CIDR_OVERLAPPING=false
export TEMPLATE_FILE=cluster-templates.yaml.tpl
```

Finally, it is possible to launch the tests:

```
make e2e
```


22.4.2 Unit tests

Most unit tests can be run directly using the *ginkgo* CLI, which in turn supports the standard testing API (*go test*, IDE features, ...). The only requirement is the *controller-runtime envtest* environment, which can be installed through *setup-envtest*:

```
go install sigs.k8s.io/controller-runtime/tools/setup-envtest@latest
setup-envtest use 1.25.x!
```

To enable the downloaded envtest, you can append the following line to your `~/.bashrc` or `~/.zshrc` file:

```
source <(setup-envtest use --installed-only --print env 1.25.x)
```

Some networking tests, however, require an isolated environment. To this end, you can leverage the dedicated *liqo-test* Docker image (the Dockerfile is available in *build/liqo-test*):

```
# Build the liqo-test Docker image
make test-container

# Run all unit tests, and retrieve coverage
make unit

# Run the tests for a specific package.
# Note, the package path must start with ./ to avoid the "package ... is not in GOROOT_
↪error".
make unit PACKAGE_PATH=<package_path>
```

Debugging unit tests

When executing the unit tests from the *liqo-test* container, it is possible to use Delve to perform remote debugging:

1. Start the *liqo-test* container with an idle entry point, exposing a port of choice (e.g. 2345):

```
docker run --name=liqo-test -d -p 2345:2345 --mount type=bind,src=$(pwd),dst=/go/
↪src/liqo \
  --privileged=true --workdir /go/src/liqo --entrypoint="" liqo-test tail -f /dev/
↪null
```

2. Open a shell inside the *liqo-test* container, and install Delve:

```
docker exec -it liqo-test bash
go install github.com/go-delve/delve/cmd/dlv@latest
```

3. Run a specific test inside the container:

```
dlv test --headless --listen=:2345 --api-version=2 \
  --accept-multiclient ./path/to/test/directory
```

4. From the host, connect to *localhost:2345* with your remote debugging client of choice (e.g. *GoLand*), and enjoy!