CLOUD DAY 2024

improove

Milano, Nov 20

Multicluster made easy: un esempio cloud-native di Federated Learning per l'AI

Claudio Lorina

Software Engineer @ ArubaKube

Giuseppe Zangari

Business & Product Lead

Kudos

CLOUD DAY 2024

Milano, Nov 20

improove

The world is becoming multi cloud

aruba.it KUBE

The world is becoming multi cloud

aruba.it KUBE

AVG Global Multi-cloud management Market

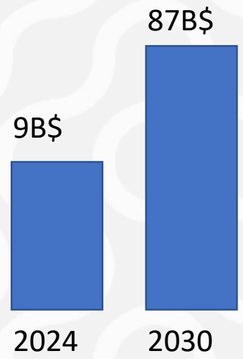

CAGR 27%

Sources

- https://www.grandviewresearch.com/industry-analysis/multi-cloud-management-market-report
- https://www.fortunebusinessinsights.com/multi-cloud-management-market-108886
- https://www.linkedin.com/pulse/global-multi-cloud-management-platform-market-trend-2032-ktppf/

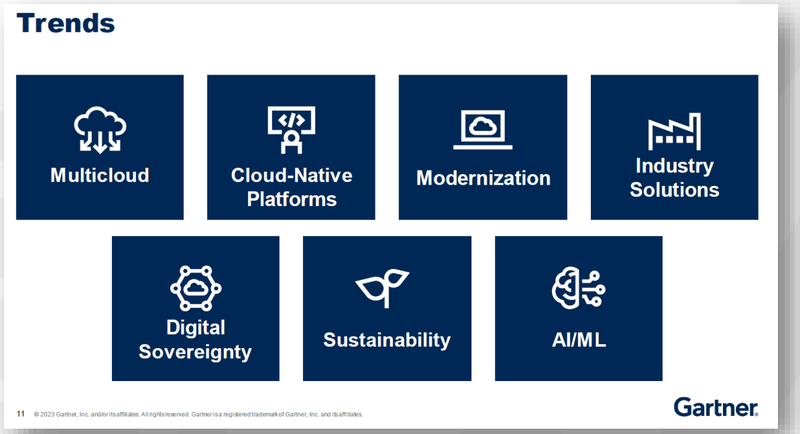

Press Release

98% of Enterprises Using Public Cloud Have Adopted a Multicloud Infrastructure Provider Strategy

New research found that IT leaders use multiple cloud providers to benefit from reduced costs and more control over their data

Austin, Texas—February 9, 2023

Why Bare Metal is Making a Comeback in a Cloud-Driven World

15 Nov 2024 by Datacenters.com Bare Metal

Free traffic inside and outside cloud

AWS News Blog

Free data transfer out to internet when moving out of AWS

by Sébastien Stormacq | on 05 MAR 2024 | in Announcements, Networking & Content Delivery, News, Price Reduction | Permalink | Comments | Share

NOW AVAILABLE

Now available: Free data transfer out to internet when leaving Azure

Published date: March 13, 2024

Networking

Cloud switching just got easier: Removing data transfer fees when moving off Google Cloud

January 12, 2024

Key resons for adopting a multi cloud solution

aruba.it KUBE

Avoiding Lock-in

Cost Efficiency

Resiliency and Redundancy

Regulatory Compliance

Access specific resources

“Let’s just add another cloud service”

Complexity: It doesn’t just grow—it compounds.

- **Interdependencies Multiply**
  One change, countless effects.
- **Unpredictability Surges**
  New layers, new surprises.
- **Duplications Create Drift**
  Aligned in theory, misaligned in practice.

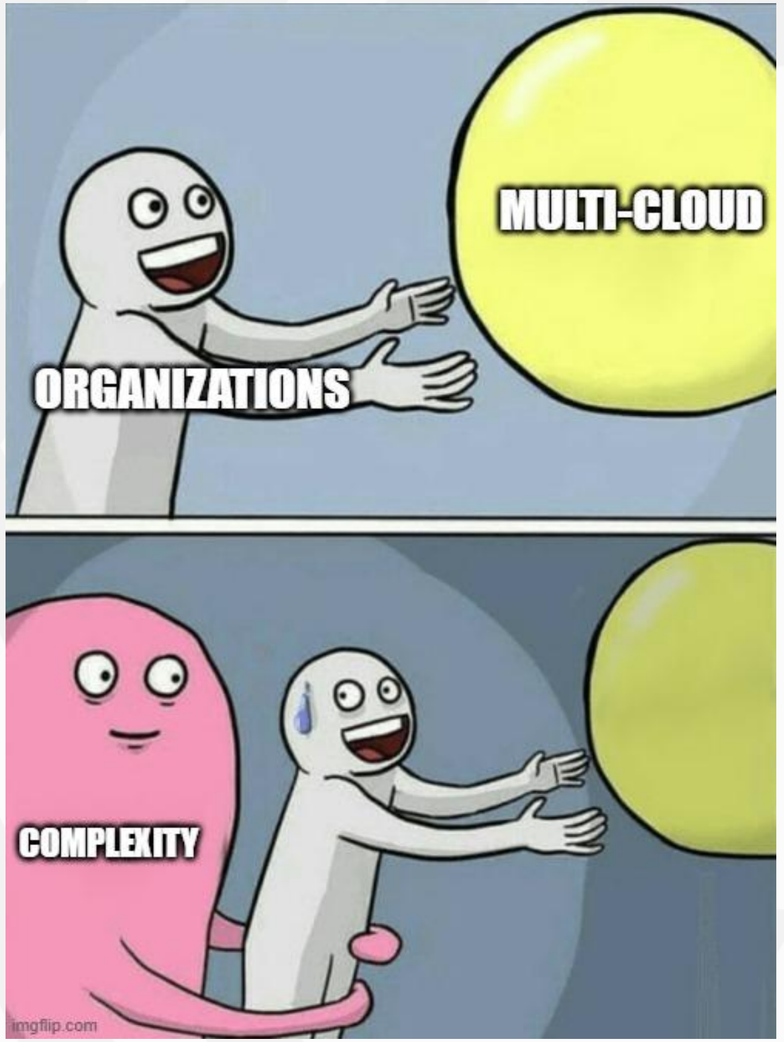

Desired scenario: More clouds for everyone

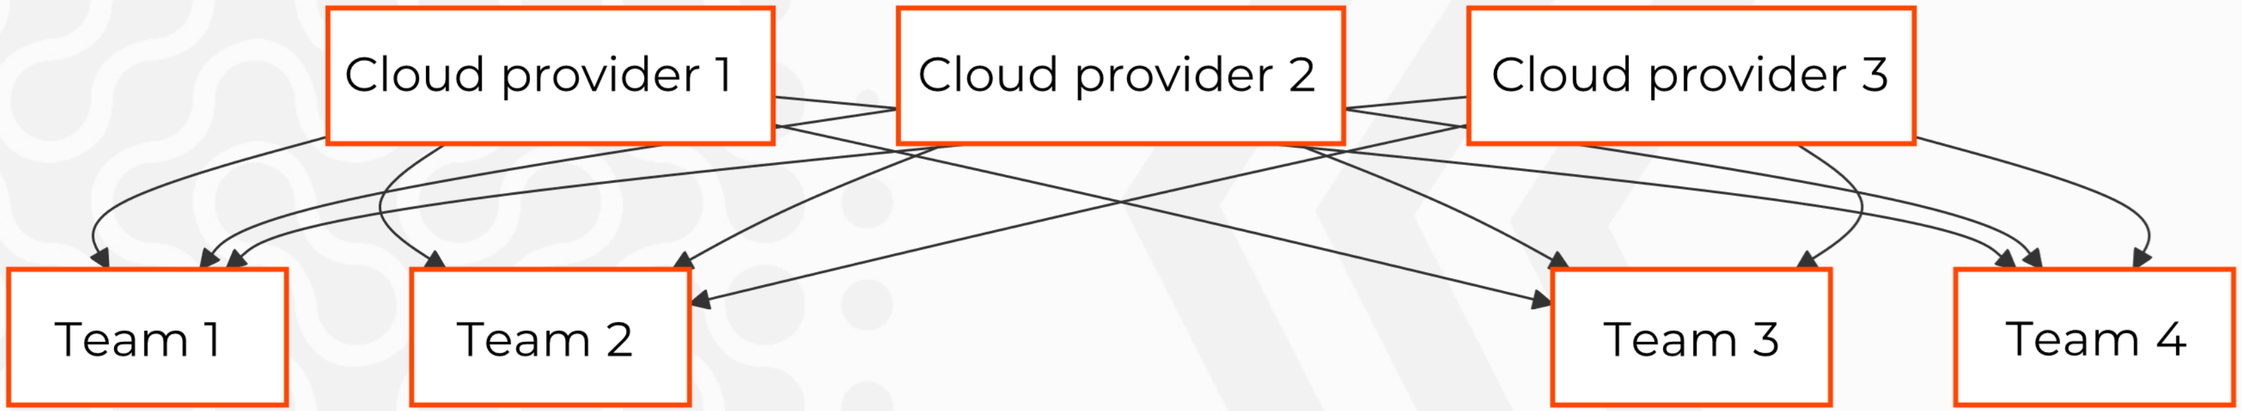

Enabling all teams to access diverse cloud resources ignites innovation and enhances product capabilities.

Actual Scenario: Siloed Access to Cloud

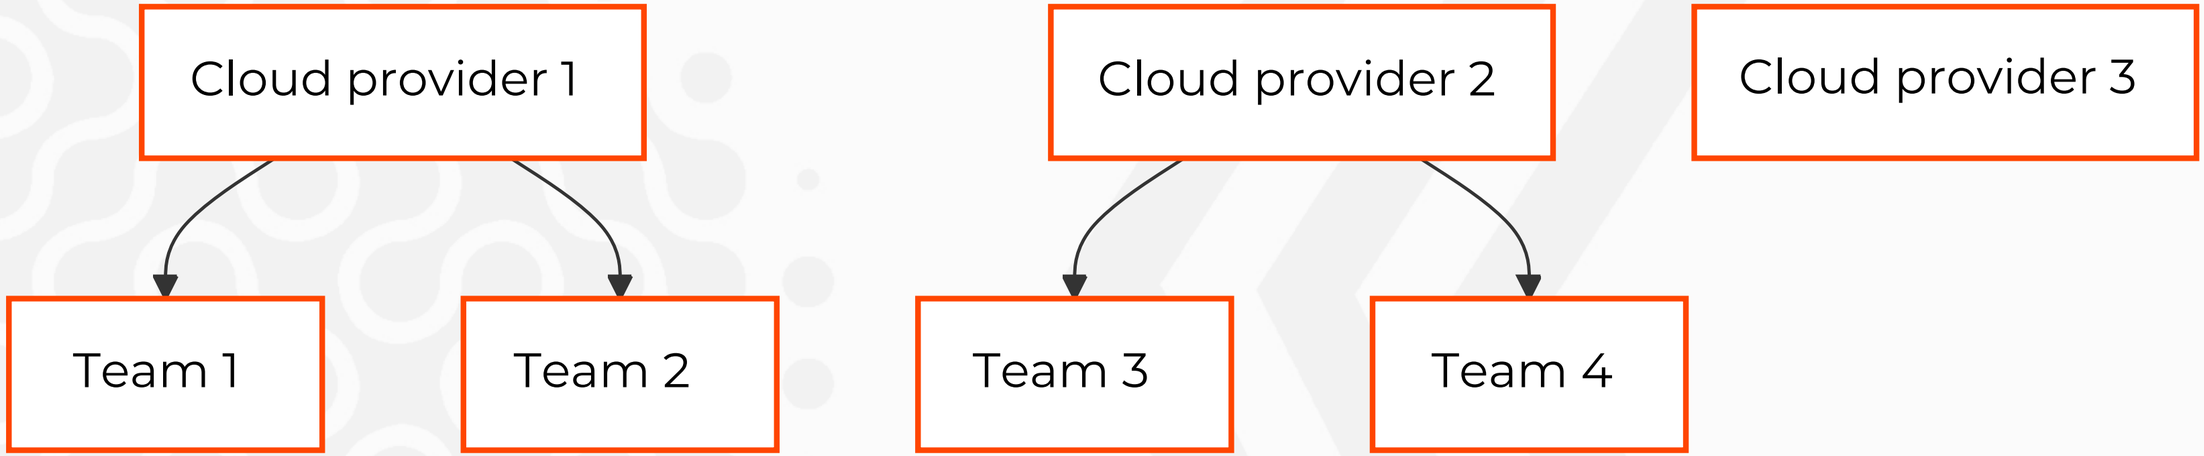

Still billing, no one uses it!

Teams limit themself to specific cloud providers, resulting in siloed access and missed opportunities for collaboration and resource sharing

Conway Law and Cognitive Load

"Any organization that designs a system will produce a design whose structure is a copy of the organization's communication structure."
Melvin Conway

Cognitive load is the amount of mental effort required to process information and perform tasks, influenced by the complexity of the work and the effectiveness of the surrounding systems and processes

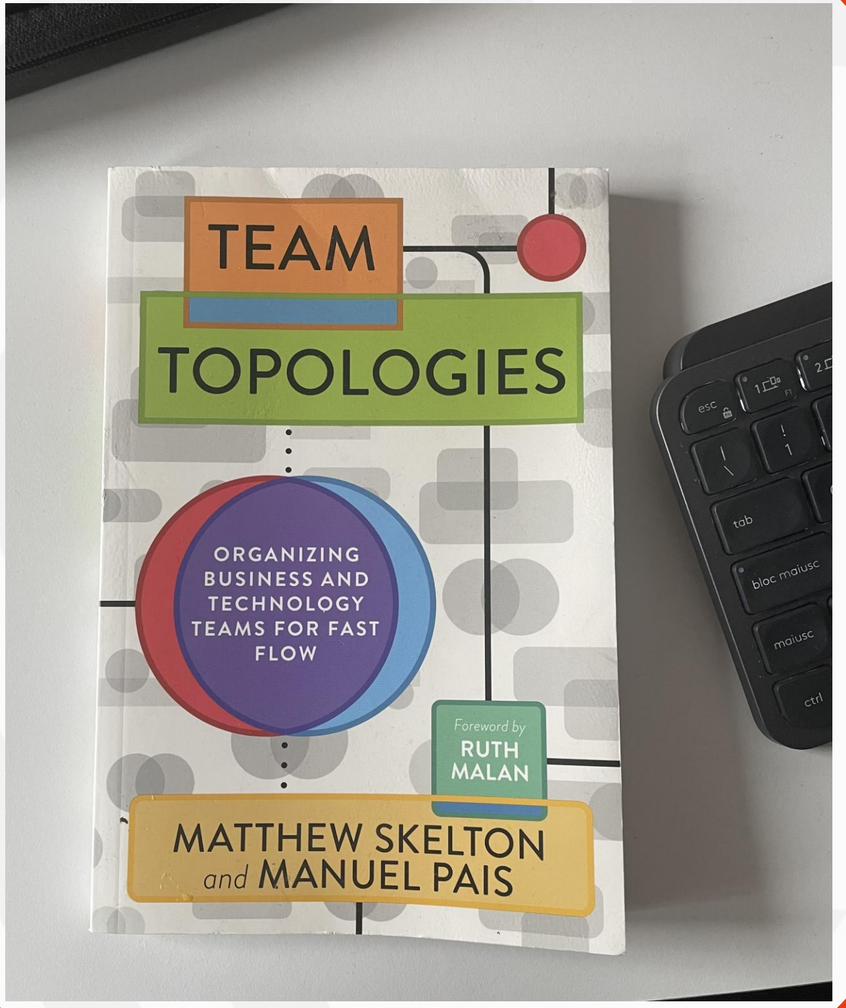

Overall complexity has 3 dimensions

Business Liability
- Slows down time-to-market and innovaiton
- Diverts resources from strategic goals

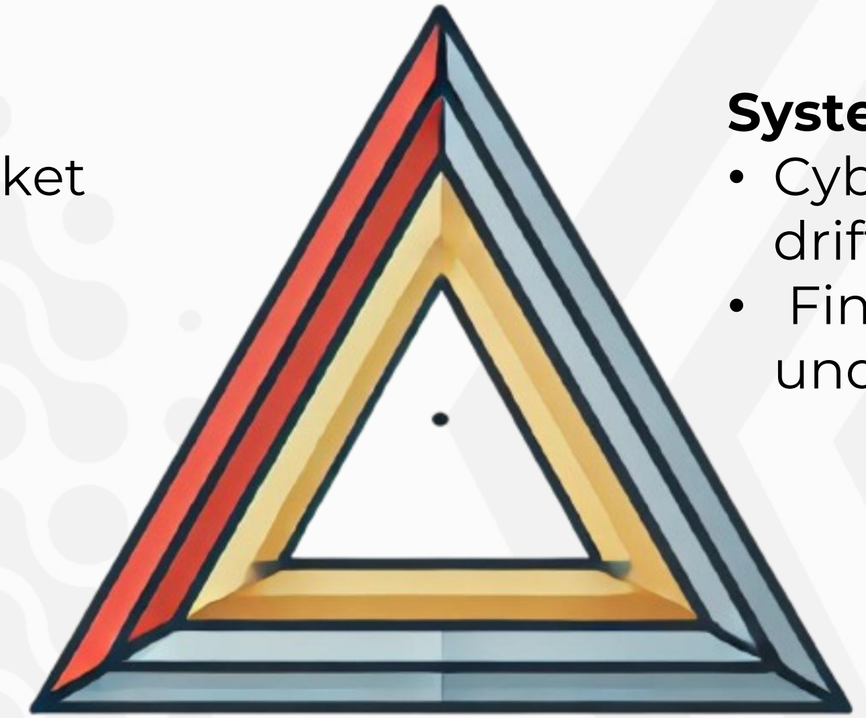

Systemic Risks
- Cybersecurity vulnerabilities and drift-related risks
- Financial risks from underutilized resources

Organizational Load
- Silos inhibit innovation and reduce cross-team alignment
- Increased cognitive load drains focus and productivity

Mastering Complexity

Adopt the right tools and processes

Mind Team and Organization Dynamics

Design and Plan with compounding effect in mind

The problem of data

Nowadays **data** has a huge strategic and economic value.

Its **potential is not always fully exploited**:

Confidential data (GDPR)

Fragmentation

Machine learning

Data is **moved** and **collected into a central dataset**

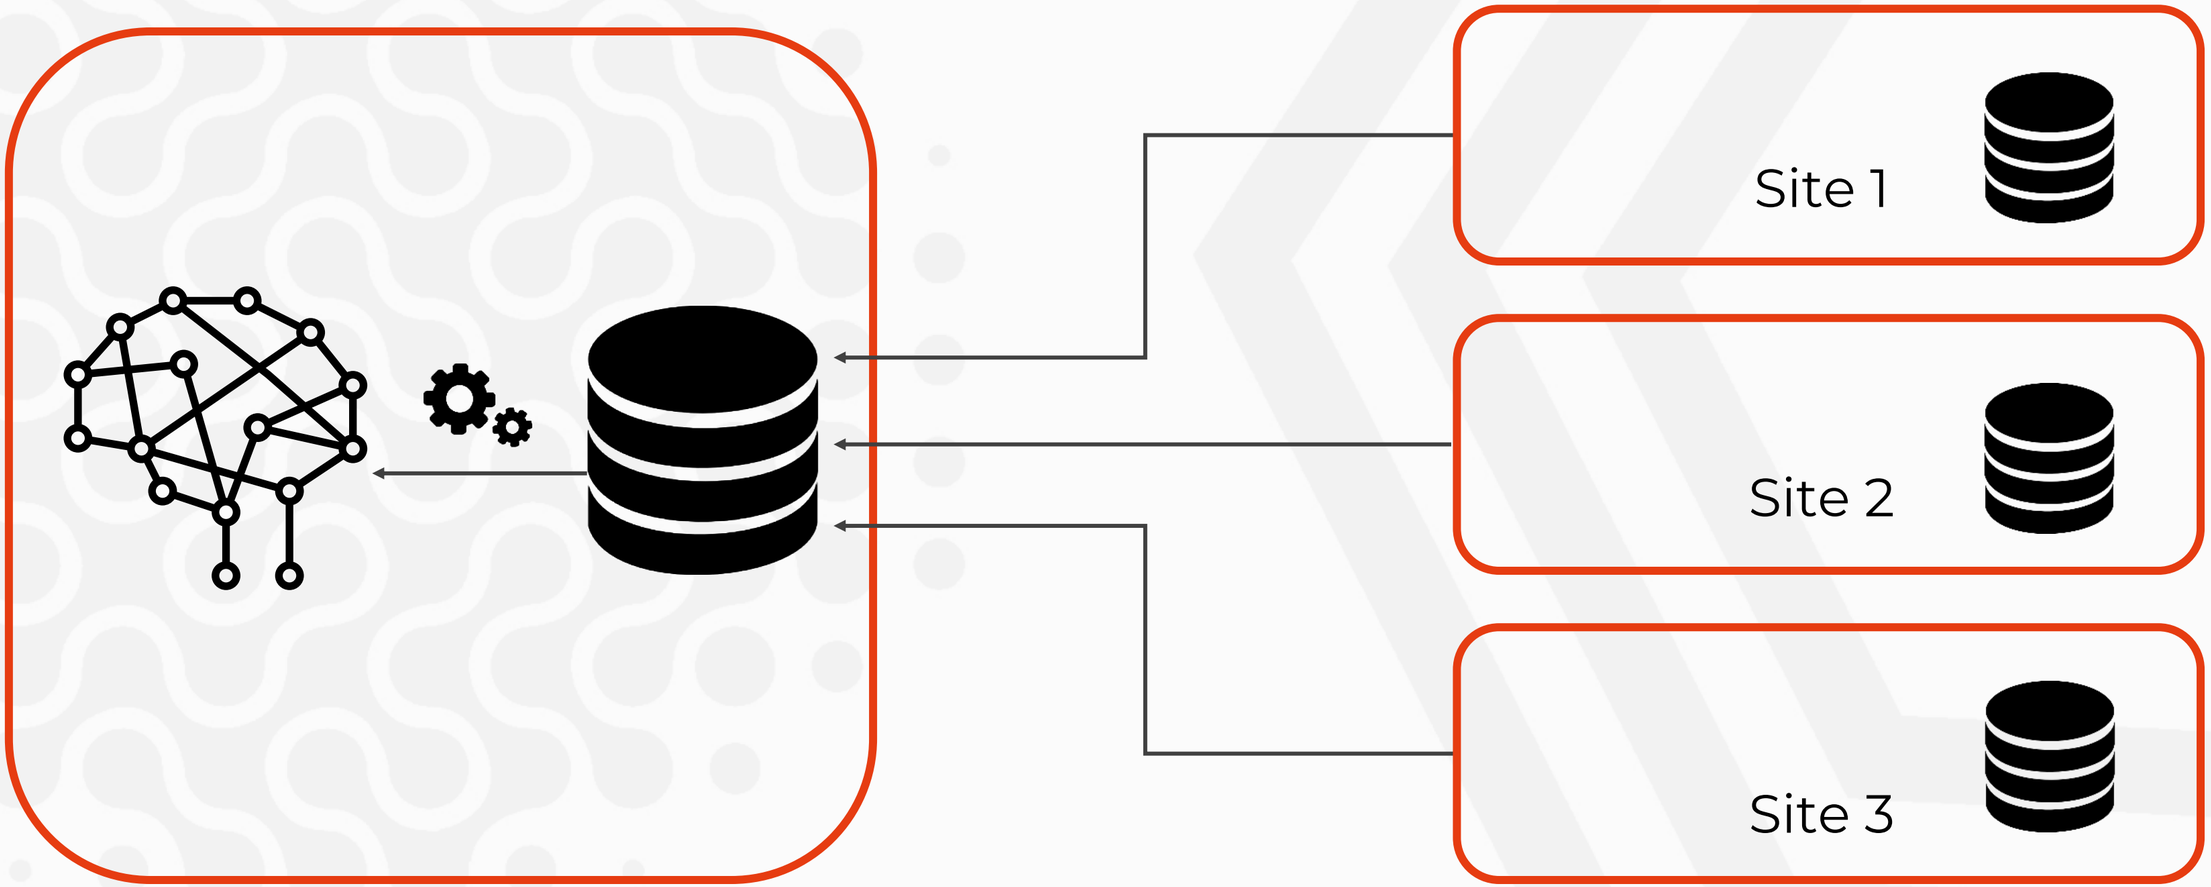

Federated learning

Data is **not** moved, **computation is moved to data**

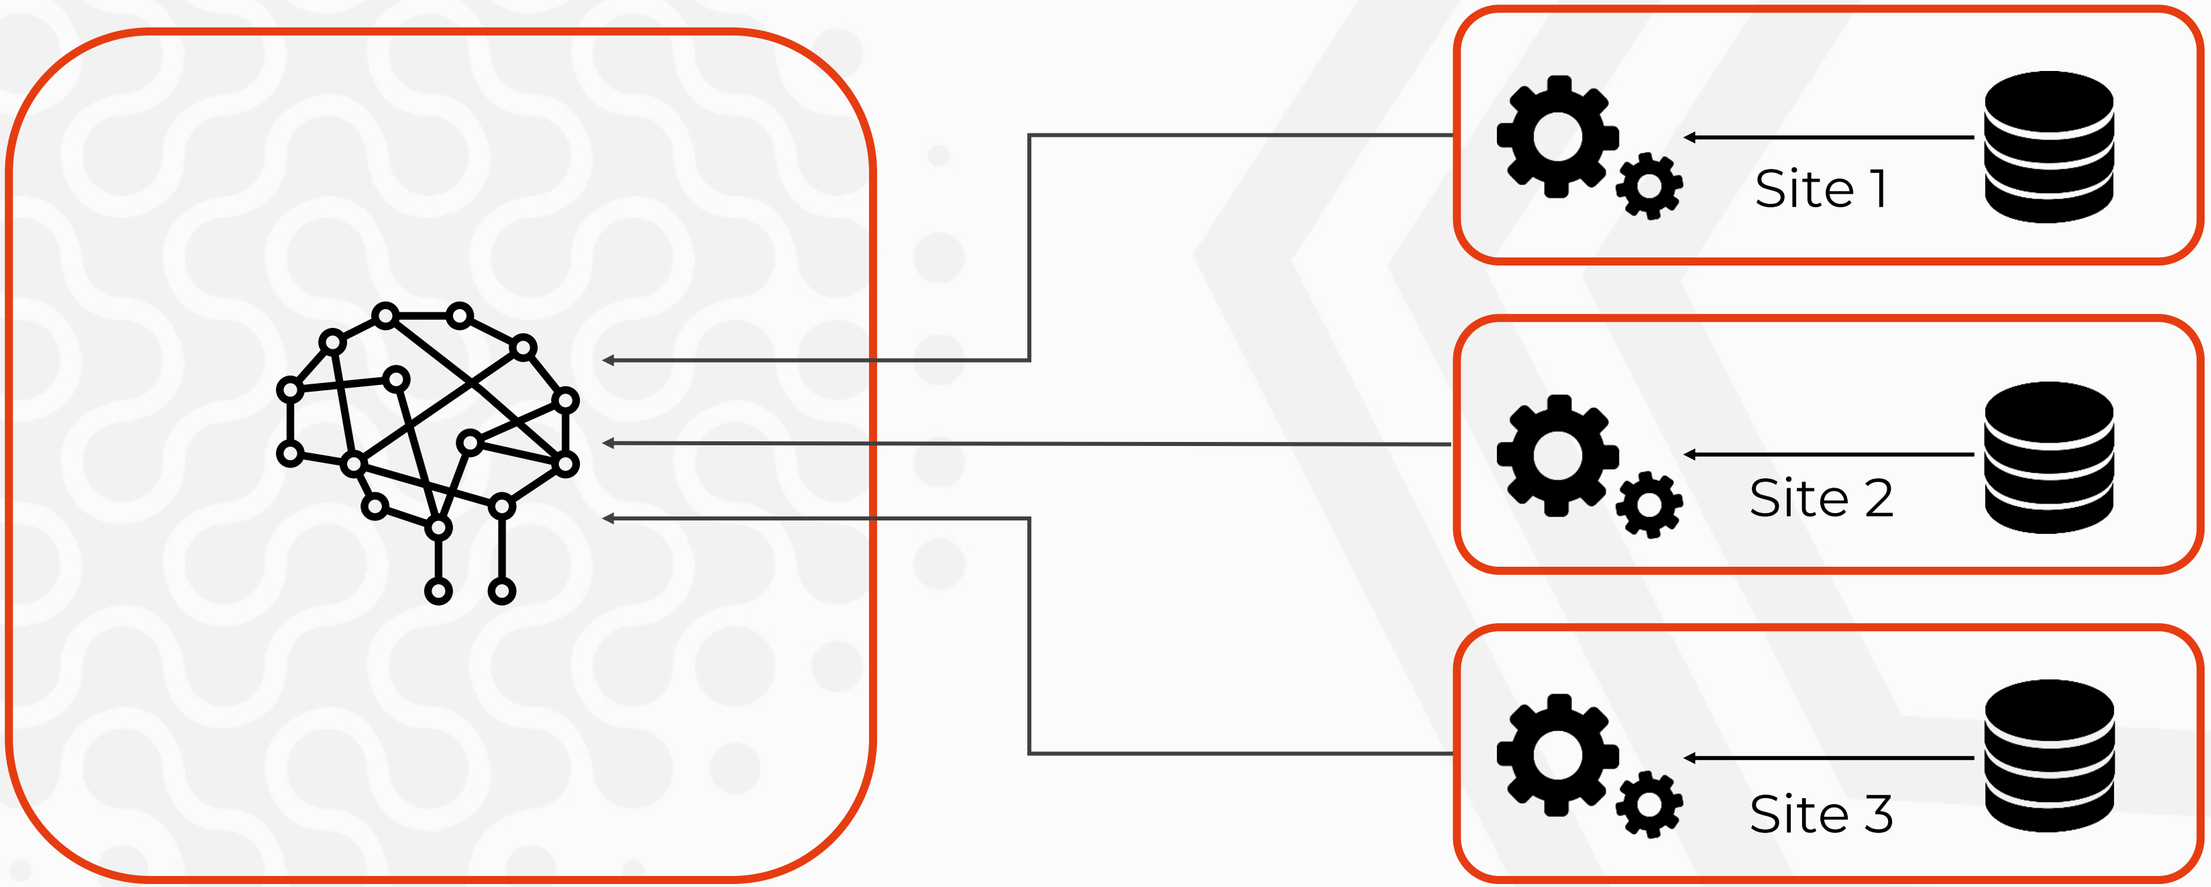

Federated learning scenario

Each hospital has its own data about patients

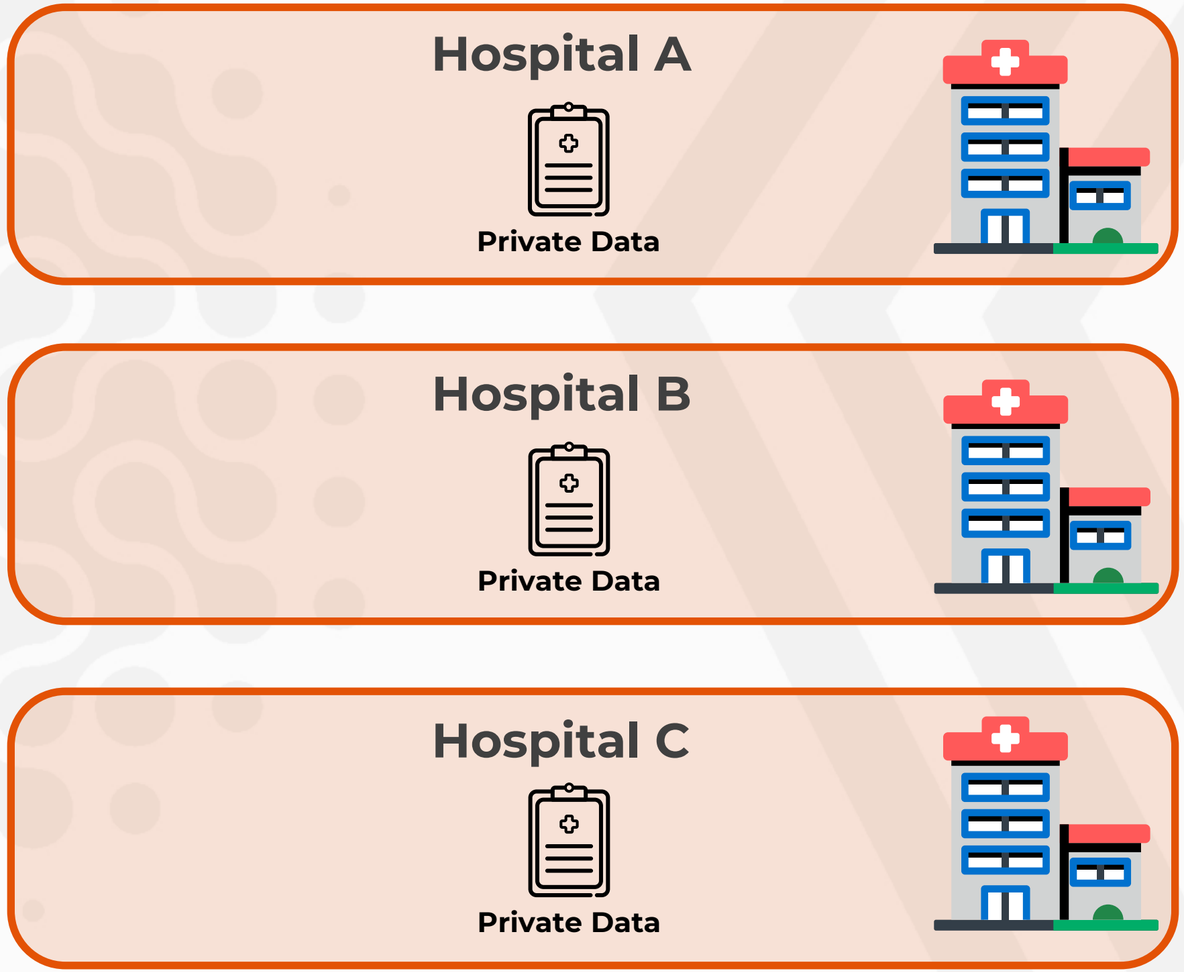

Federated learning scenario

We might use the data to train local models

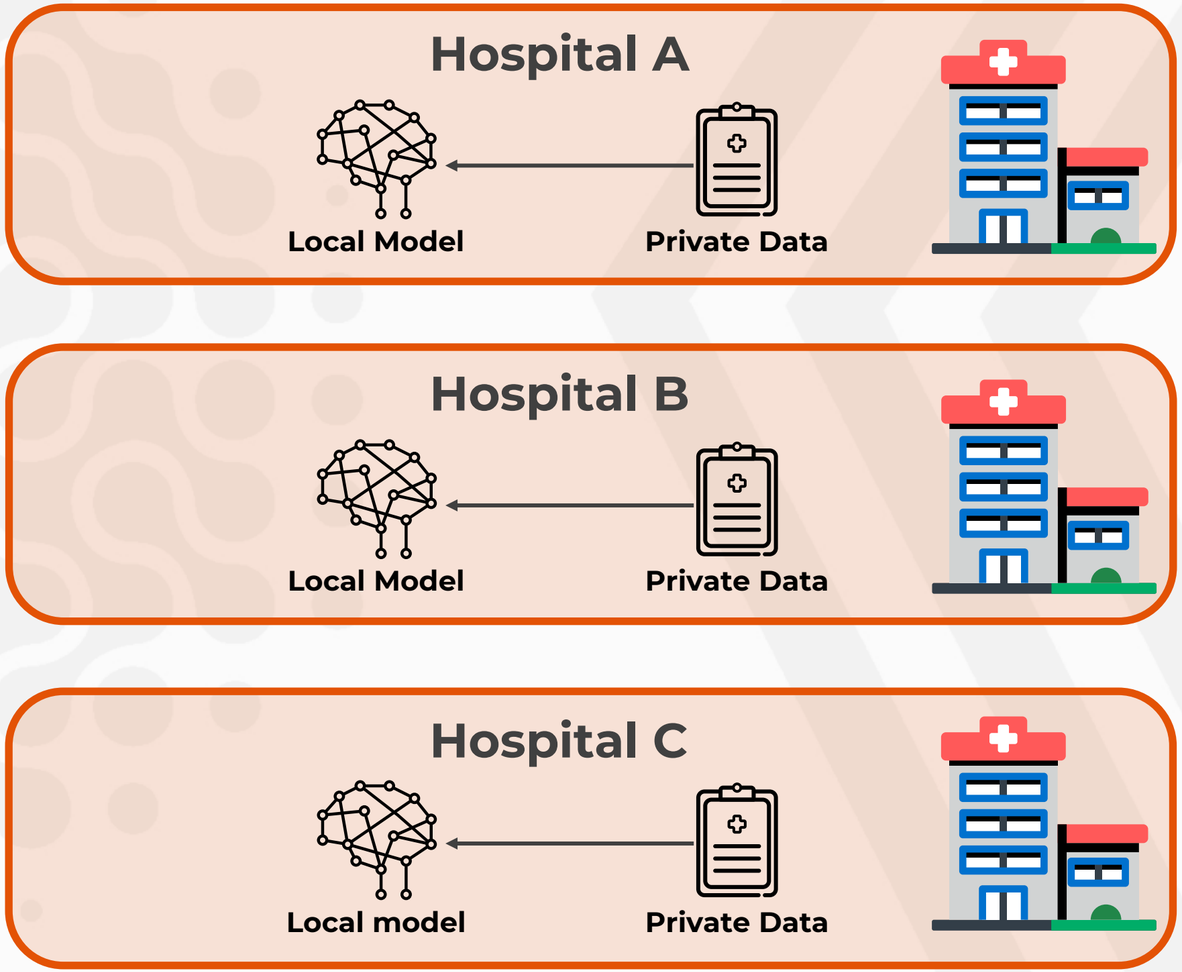

Federated learning scenario

With federated learning we **aggregate the local models, obtaining a global model with higher accuracy**

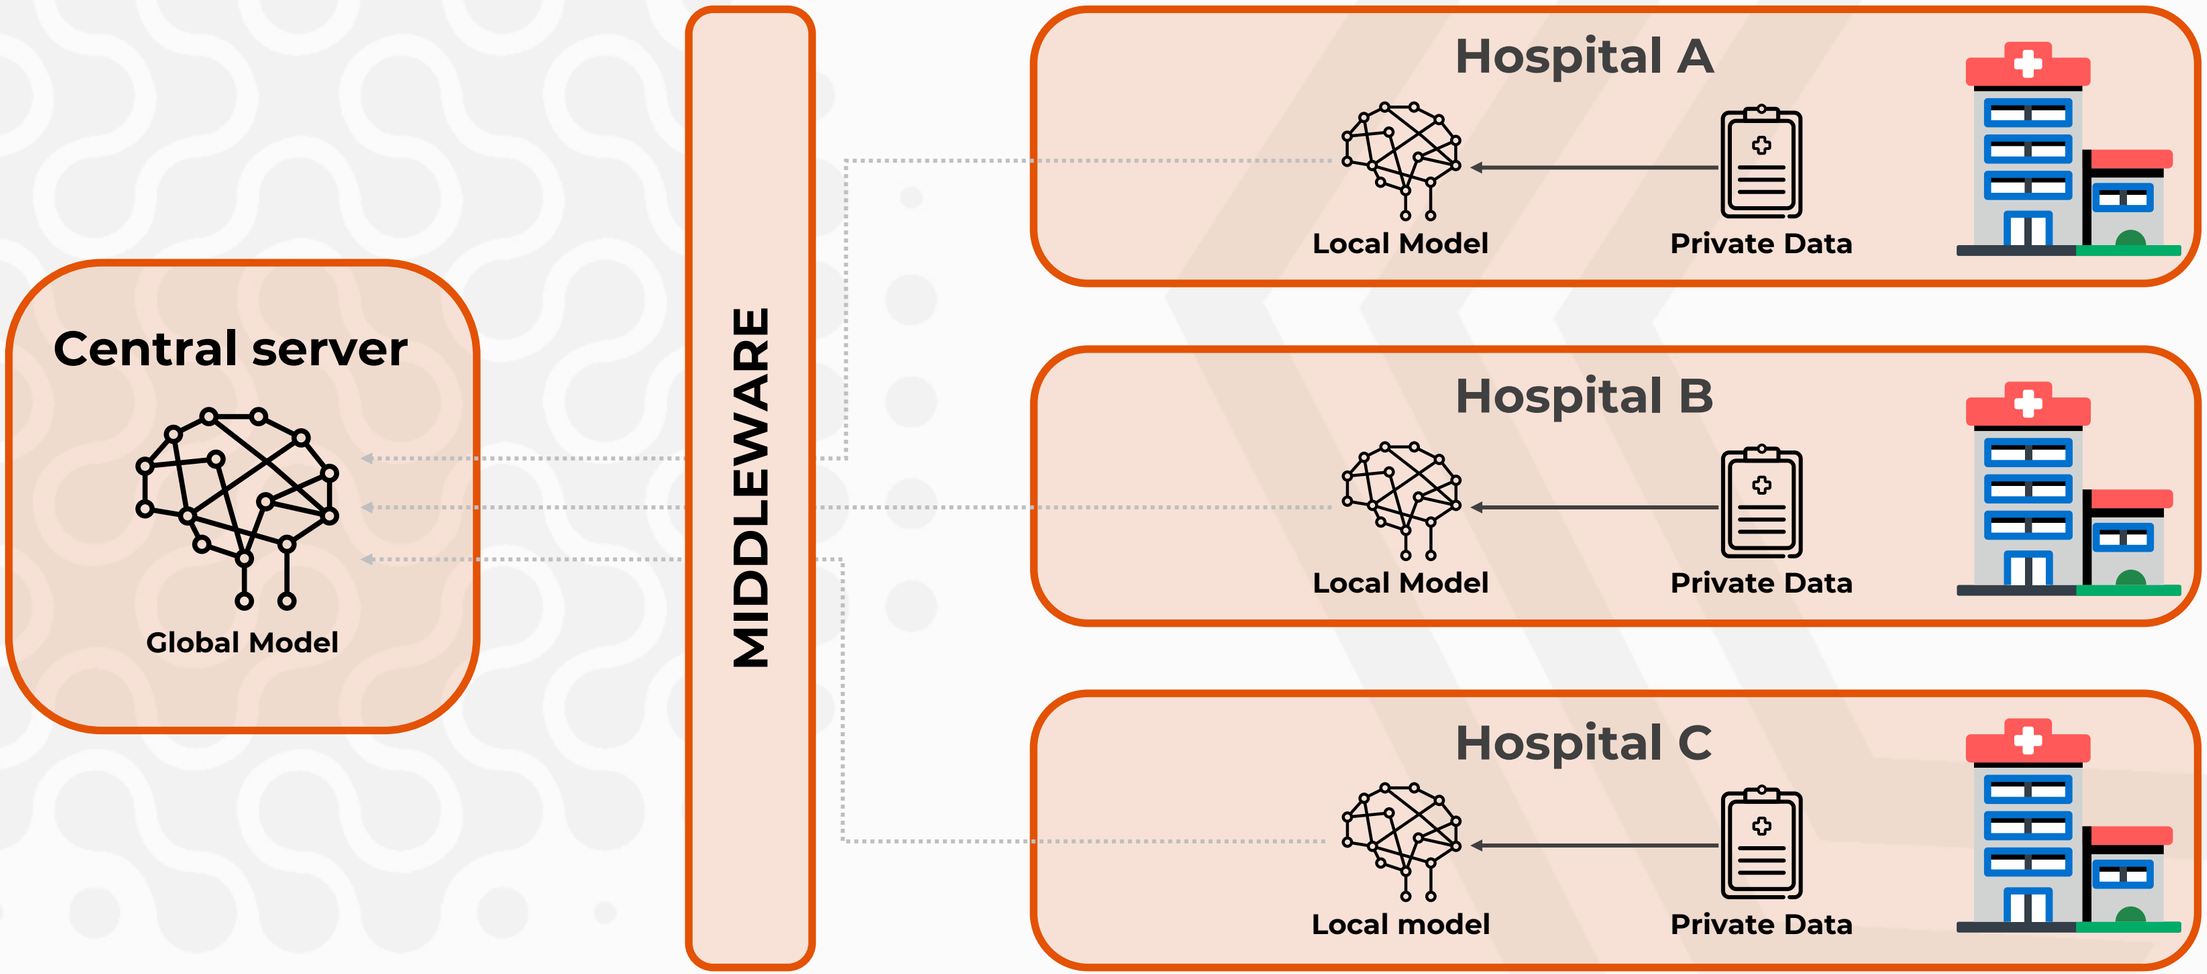

Implementation of a solution: the challenge

aruba.it KUBE

We need a solution coordinating all the involved **entities, application and infrastructure**.

Expected Result

Implementation of a solution: the challenge

aruba.it KUBE

When multiple things and people need to be coordinated,

sometimes, the result is not the expected one

Actual result

Liqo: an integrated tool for multicluster

Liqo is an **open-source** project enabling
dynamic and seamless Kubernetes multi-cluster topologies

An **integrated** solution allowing to:

Share resources and offload computing
the remote cluster becomes a node of your local cluster via *Virtual Kubelet*

Enable networking between peered clusters
enable networking between pods and toward services in different clusters

K8s resources available across all the peered clusters
ConfigMaps, Ingresses etc. are available in all the peered clusters

Liqo: seamless multicluster

Given a namespace on a cluster we can enable offloading with Liqo

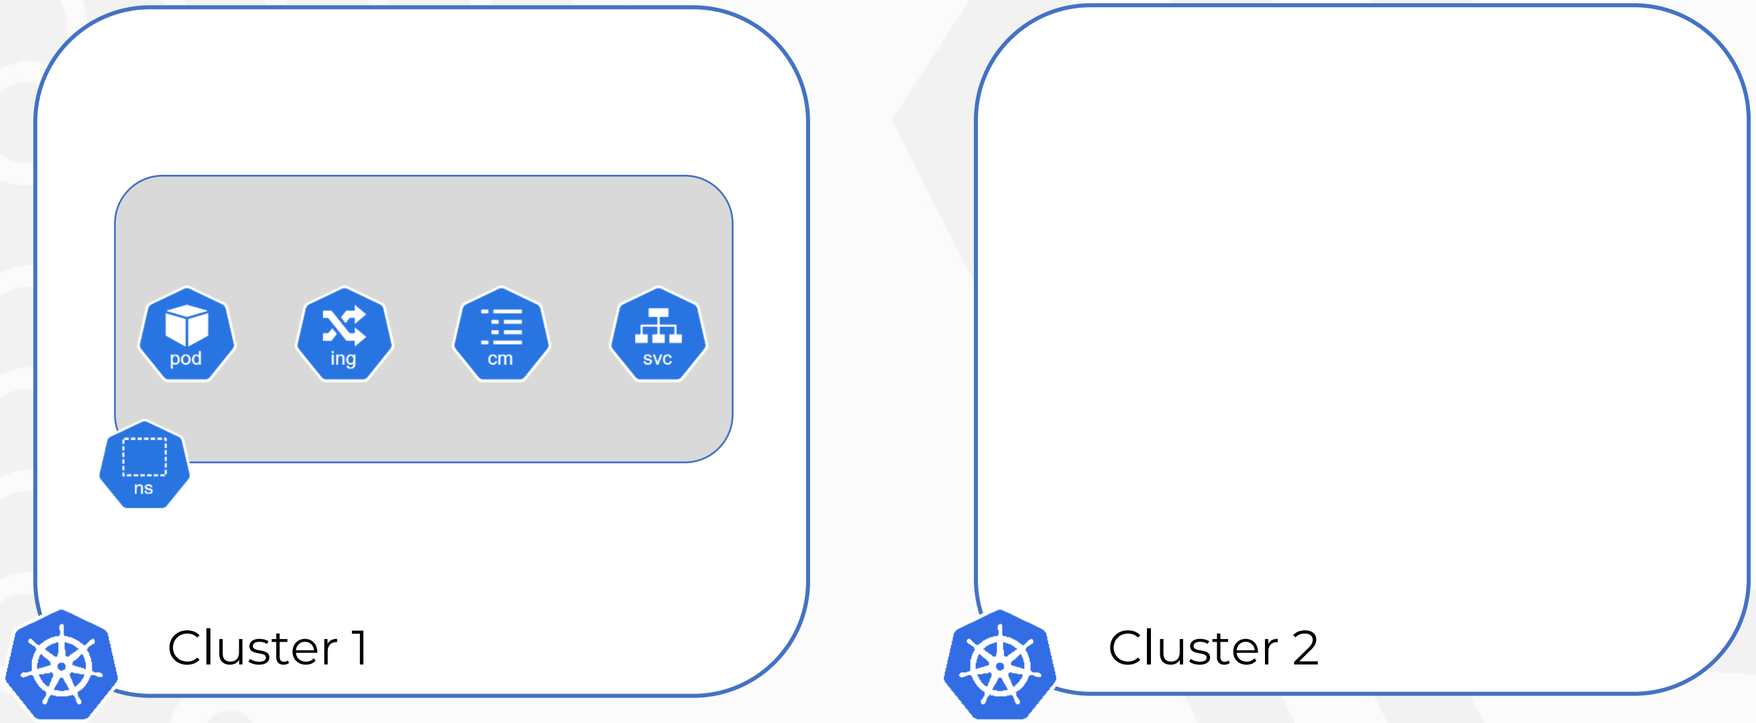

Liqo: seamless multicluster

With Liqo an “**offloaded namespace” transparently spans across multiple clusters**, **like it was a namespace in a single** Kubernetes **cluster.**

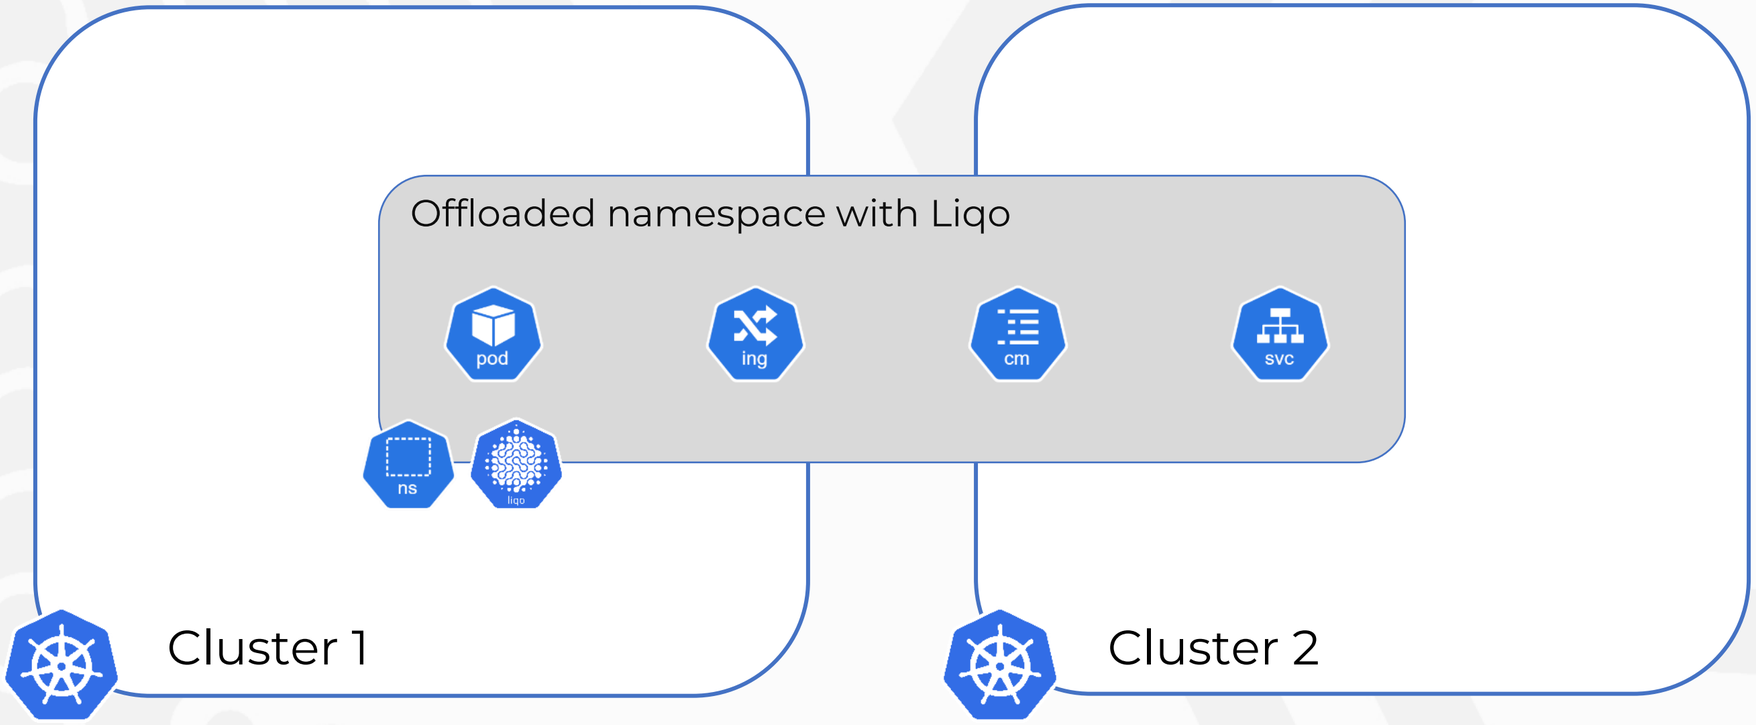

Liqo: configurability and flexibility

Highly configurable and extendable
each module can be individually configured and, if required, extended

Fully declarative configuration
easy to integrate with automation pipelines

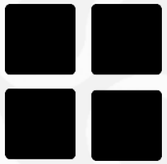

Platform agnostic
support for the major K8s distributions both bare metal or as a service.

kubeadm K3S Azure aws Google Cloud aruba CLOUD

Flower: a friendly federated AI framework

Flower is an **open-source framework** for building **federated learning systems**

Easy to use and highly customizable
It adapts to many use cases and few lines of code are required to obtain a working federated learning system.

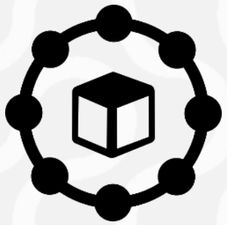

Framework and platform agnostic
Can be used with the most popular machine learning frameworks (TensorFlow, PyTorch) and it can run both in cloud or at the edge.

Federated learning with Liqo

aruba.it KUBE

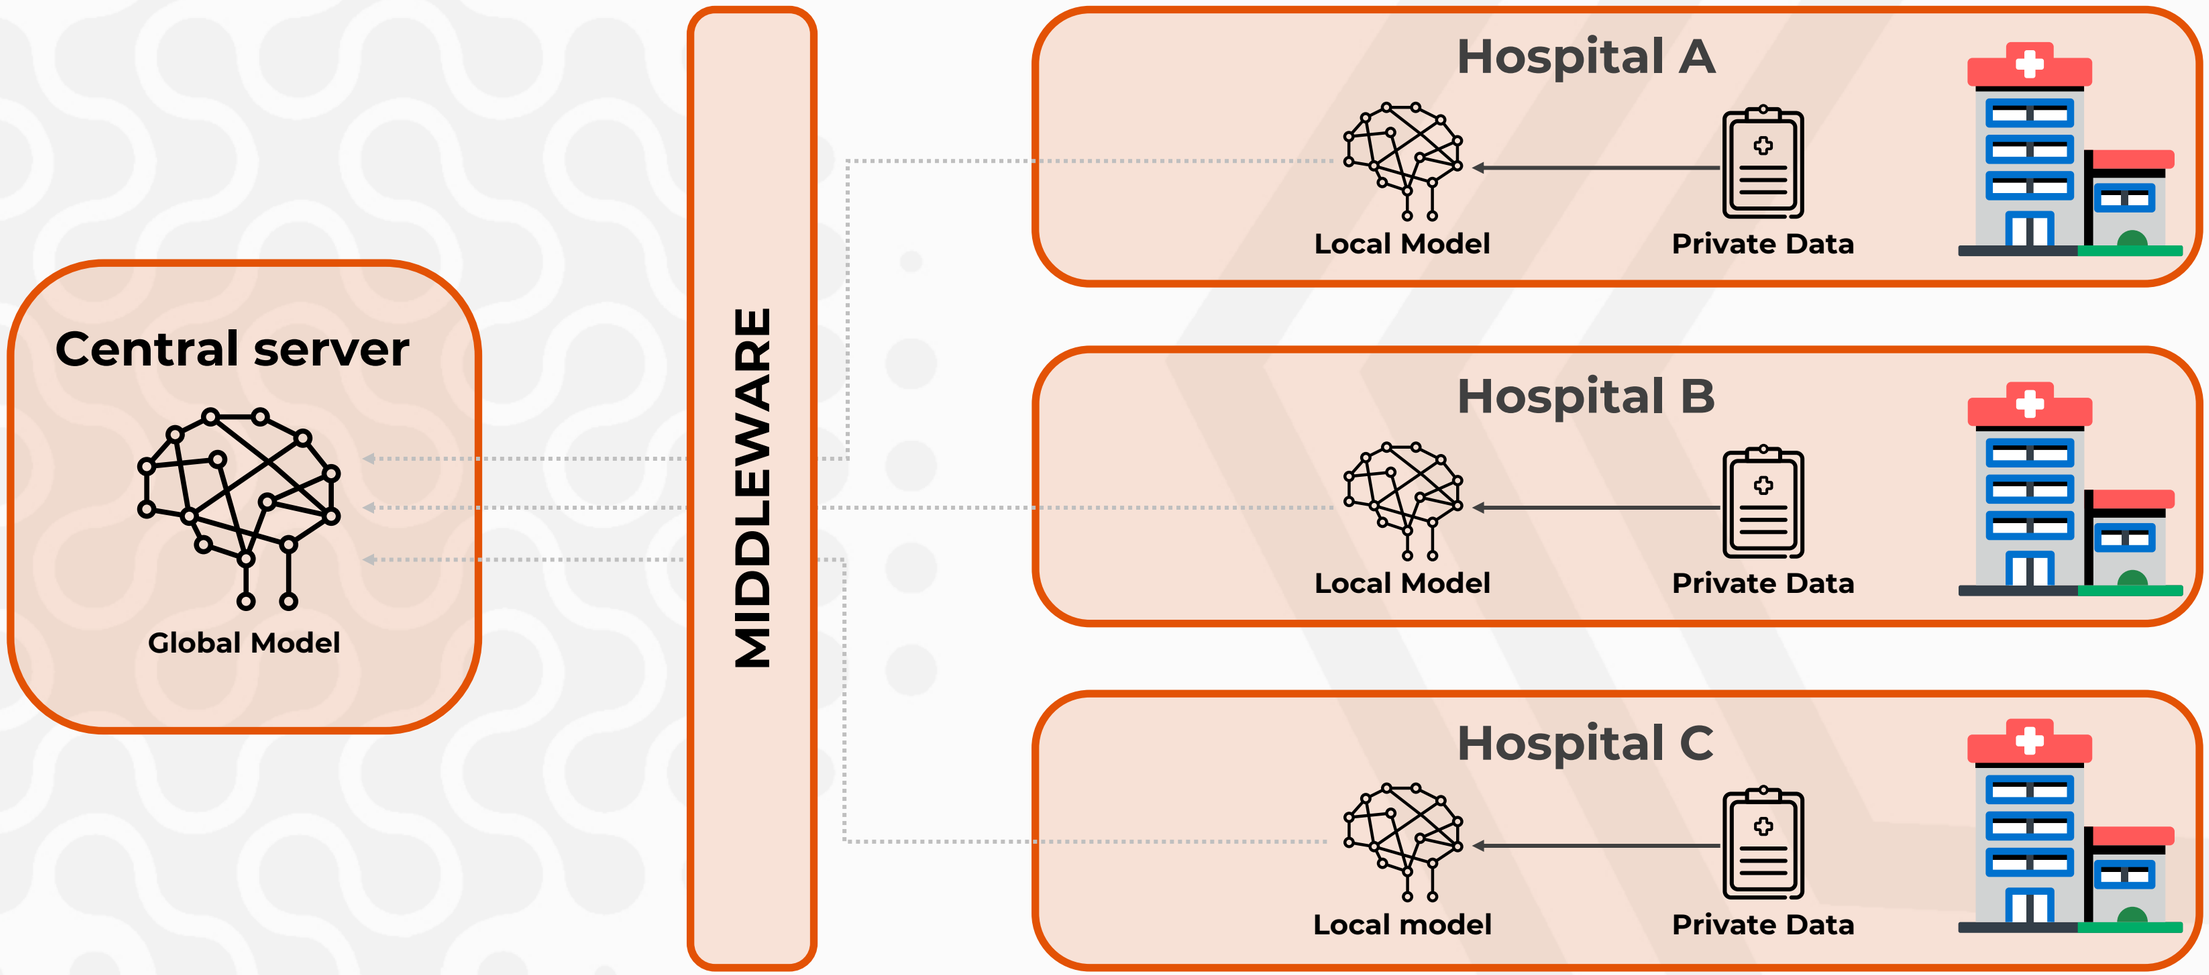

Federated learning with Liqo

Via **Liqo** we provide a **centralized control of the application and multi-cluster connectivity**

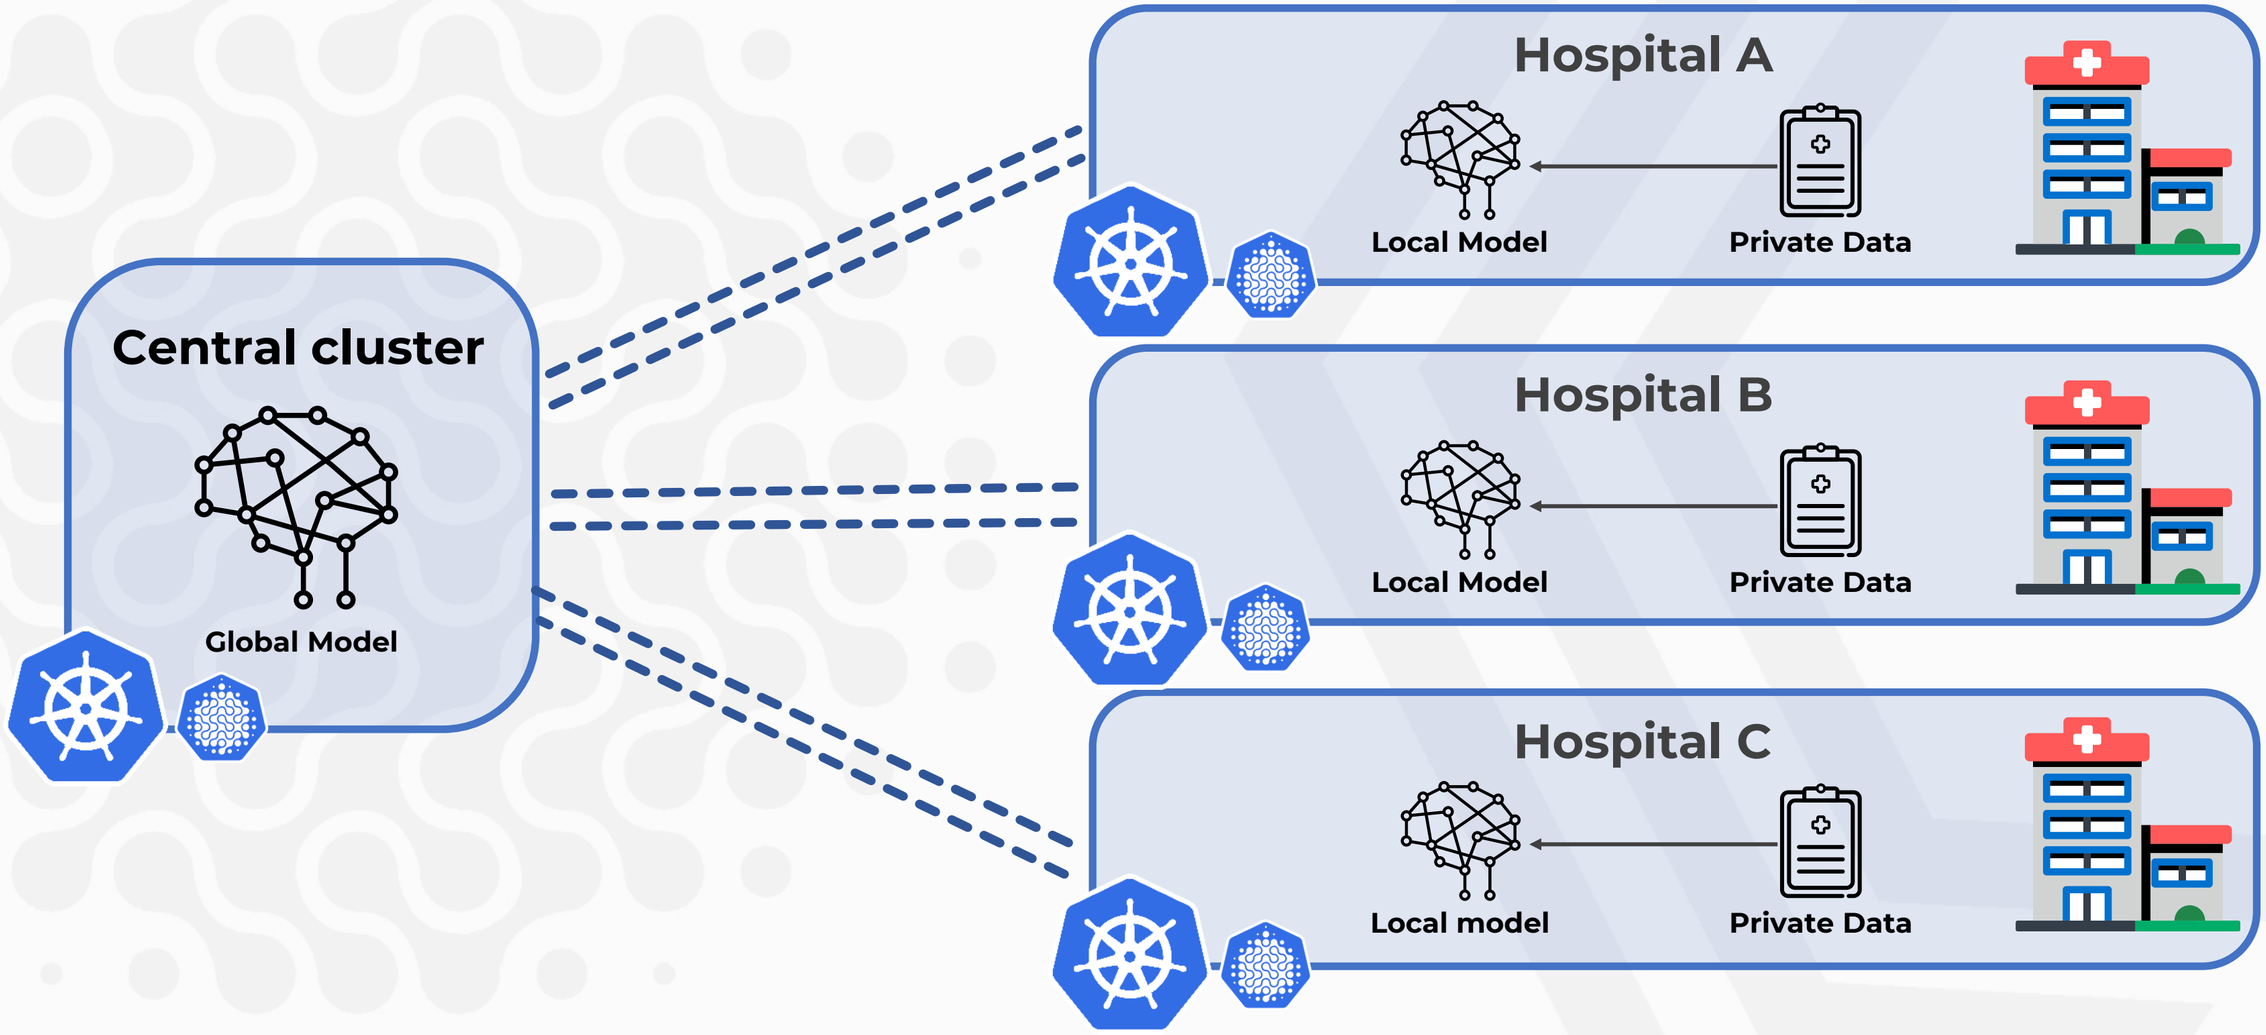

```
~ » liqoctl peer --remote-context hospital01                    (⎈ central-cluster:default)
 INFO   (local) Network configuration correctly retrieved
 INFO   (remote) Network configuration correctly retrieved
 INFO   (local) Network configuration correctly set up
 INFO   (remote) Network configuration correctly set up
 INFO   (local) Configuration applied successfully
 INFO   (remote) Configuration applied successfully
 INFO   (local) Network correctly initialized
 INFO   (remote) Network correctly initialized
 INFO   (remote) Gateway server correctly set up
 INFO   (remote) Gateway pod gw-central-cluster is ready
 INFO   (remote) Gateway server Service created successfully
 INFO   (local) Gateway client correctly set up
 INFO   (local) Gateway pod gw-hospital01 is ready
 INFO   (remote) Gateway server Secret created successfully
 INFO   (local) Public key correctly created
 INFO   (local) Gateway client Secret created successfully
 INFO   (remote) Public key correctly created
 INFO   (remote) Connection created successfully
 INFO   (local) Connection created successfully
 INFO   (local) Connection is established
 INFO   (remote) Connection is established
 INFO   (local) Tenant namespace correctly ensured
 INFO   (remote) Tenant namespace correctly ensured
 INFO   (remote) Nonce secret ensured
 INFO   (remote) Nonce generated successfully
 INFO   (remote) Nonce retrieved
 INFO   (local) Signed nonce secret ensured
 INFO   (local) Nonce is signed
 INFO   (local) Signed nonce retrieved
 INFO   (local) Tenant correctly generated
 INFO   (remote) Tenant correctly applied on provider cluster
 INFO   (remote) Tenant status is filled
 INFO   (remote) Identity correctly generated
 INFO   (local) Identity correctly applied on consumer cluster
 INFO   (local) Identity status is filled
 INFO   (local) ResourceSlice created
 INFO   (local) ResourceSlice authentication: Accepted
 INFO   (local) ResourceSlice resources: Accepted

~ »                                                            (⎈ central-cluster:default)
```

```
Every 2.0s: kubectl get nodes                  tsdi1xx-akube01.ad.aruba.it: Wed Nov 13 16:58:45 2024

NAME                           STATUS   ROLES           AGE    VERSION
central-cluster-control-plane  Ready    control-plane   5h4m   v1.30.0
central-cluster-worker         Ready    <none>          5h4m   v1.30.0
hospital01                     Ready    agent           1s     v1.30.0
```

Liqo virtual nodes

Liqo abstract the hospitals' **clusters making them nodes on the central cluster**

```
~ » kubectl get nodes
NAME                                STATUS   ROLES           AGE     VERSION
central-cluster-control-plane       Ready    control-plane   5h59m   v1.30.0
central-cluster-worker              Ready    <none>          5h59m   v1.30.0
hospital01                          Ready    agent           55m     v1.30.0
hospital02                          Ready    agent           54m     v1.30.0
hospital03                          Ready    agent           54m     v1.30.0
```

Liqo Virtual Nodes

Extend the namespace on the remote clusters

Doing namespace offloading we **extend** the *“flower-demo”* **namespace** on the clusters o the hospitals

```
~ » liqoctl offload namespace flower-demo                    (⎈ central-cluster:default)
 INFO  Offloading of namespace "flower-demo" correctly enabled
 INFO  Offloading completed successfully
```

FL with Liqo and flower: demo

aruba.it KUBE

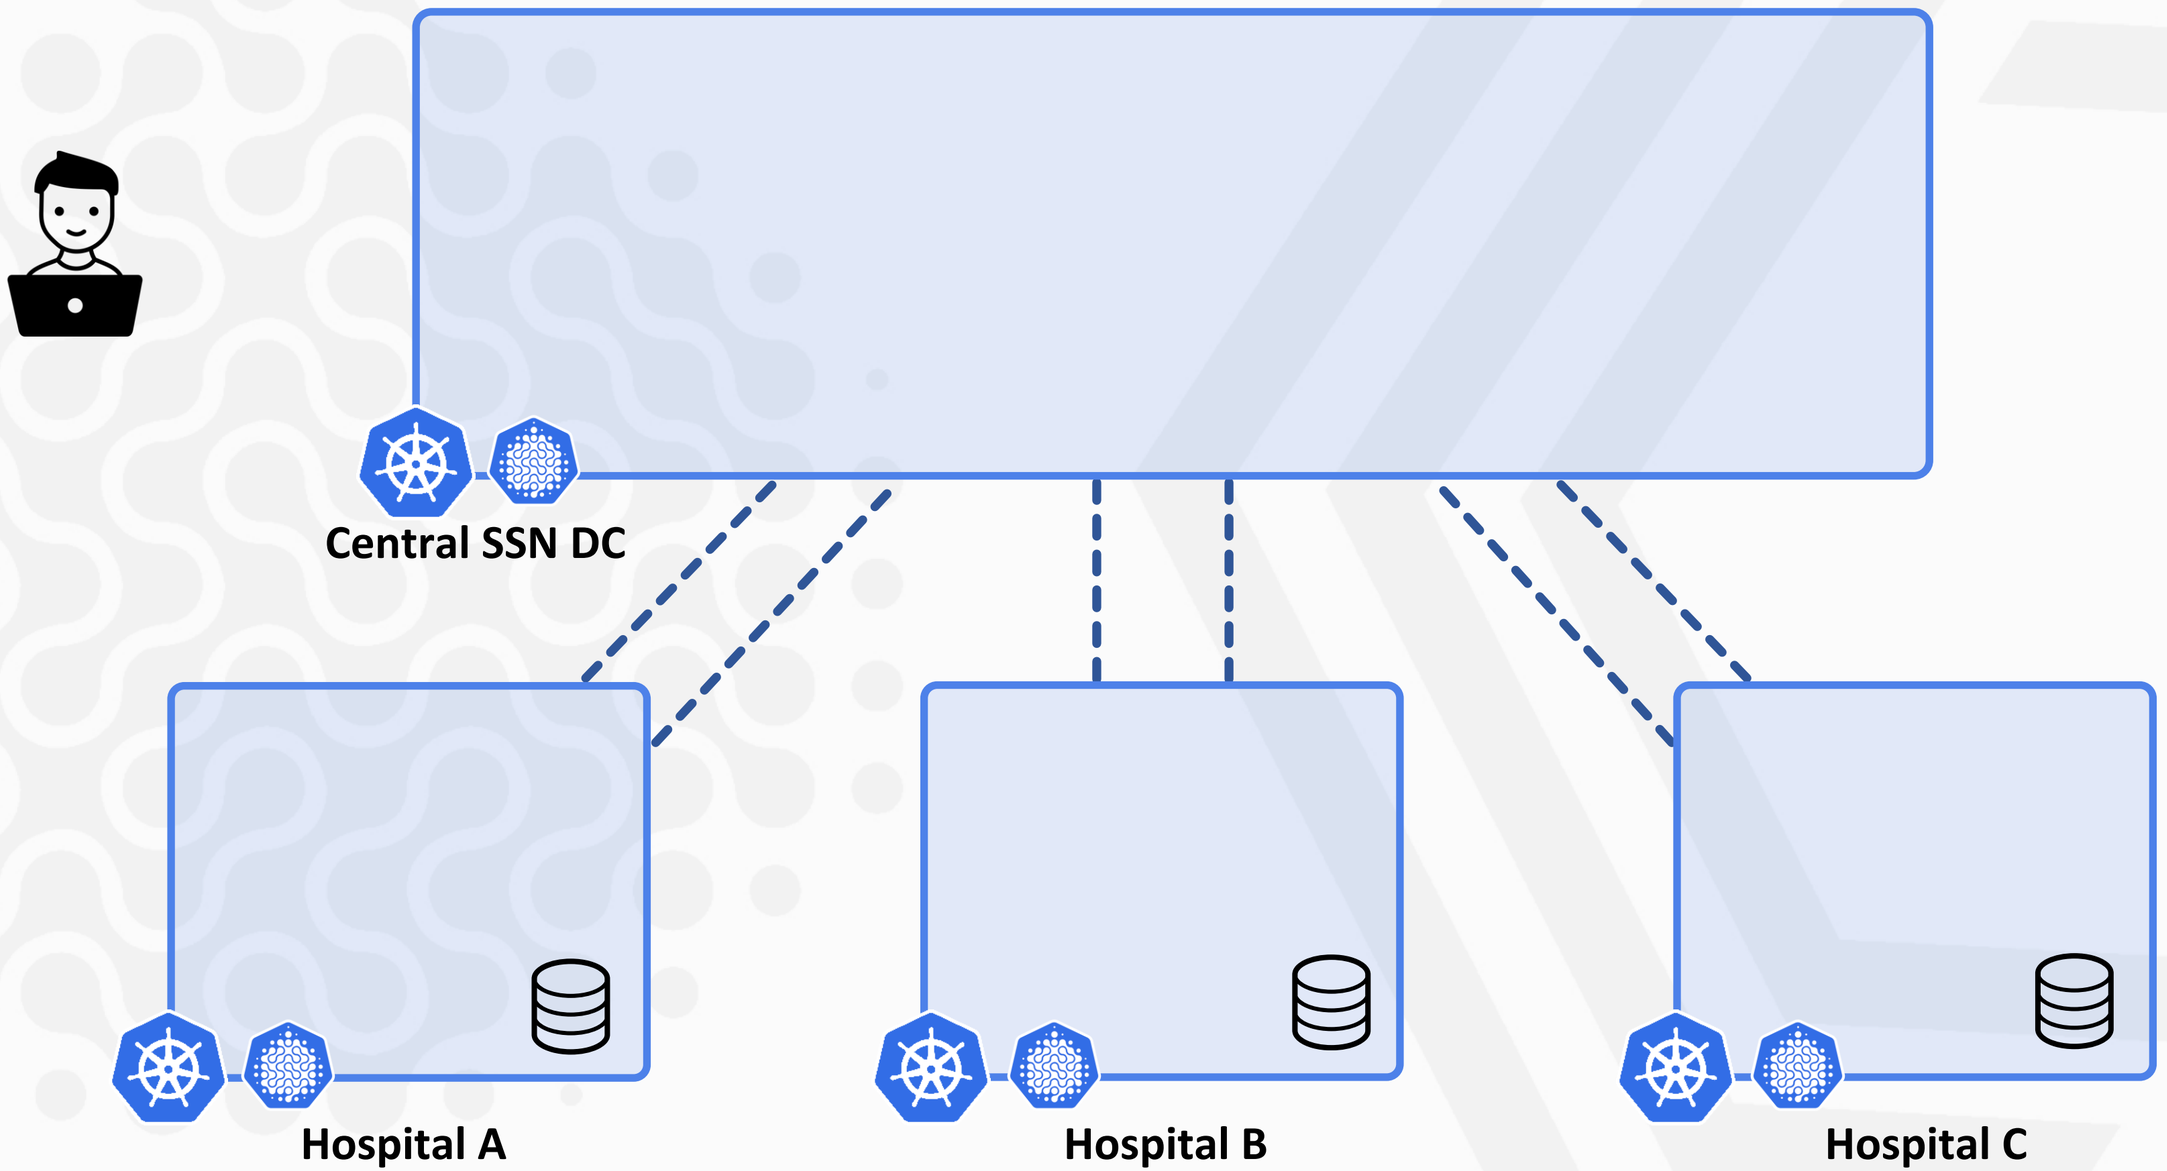

FL with Liqo and flower: demo

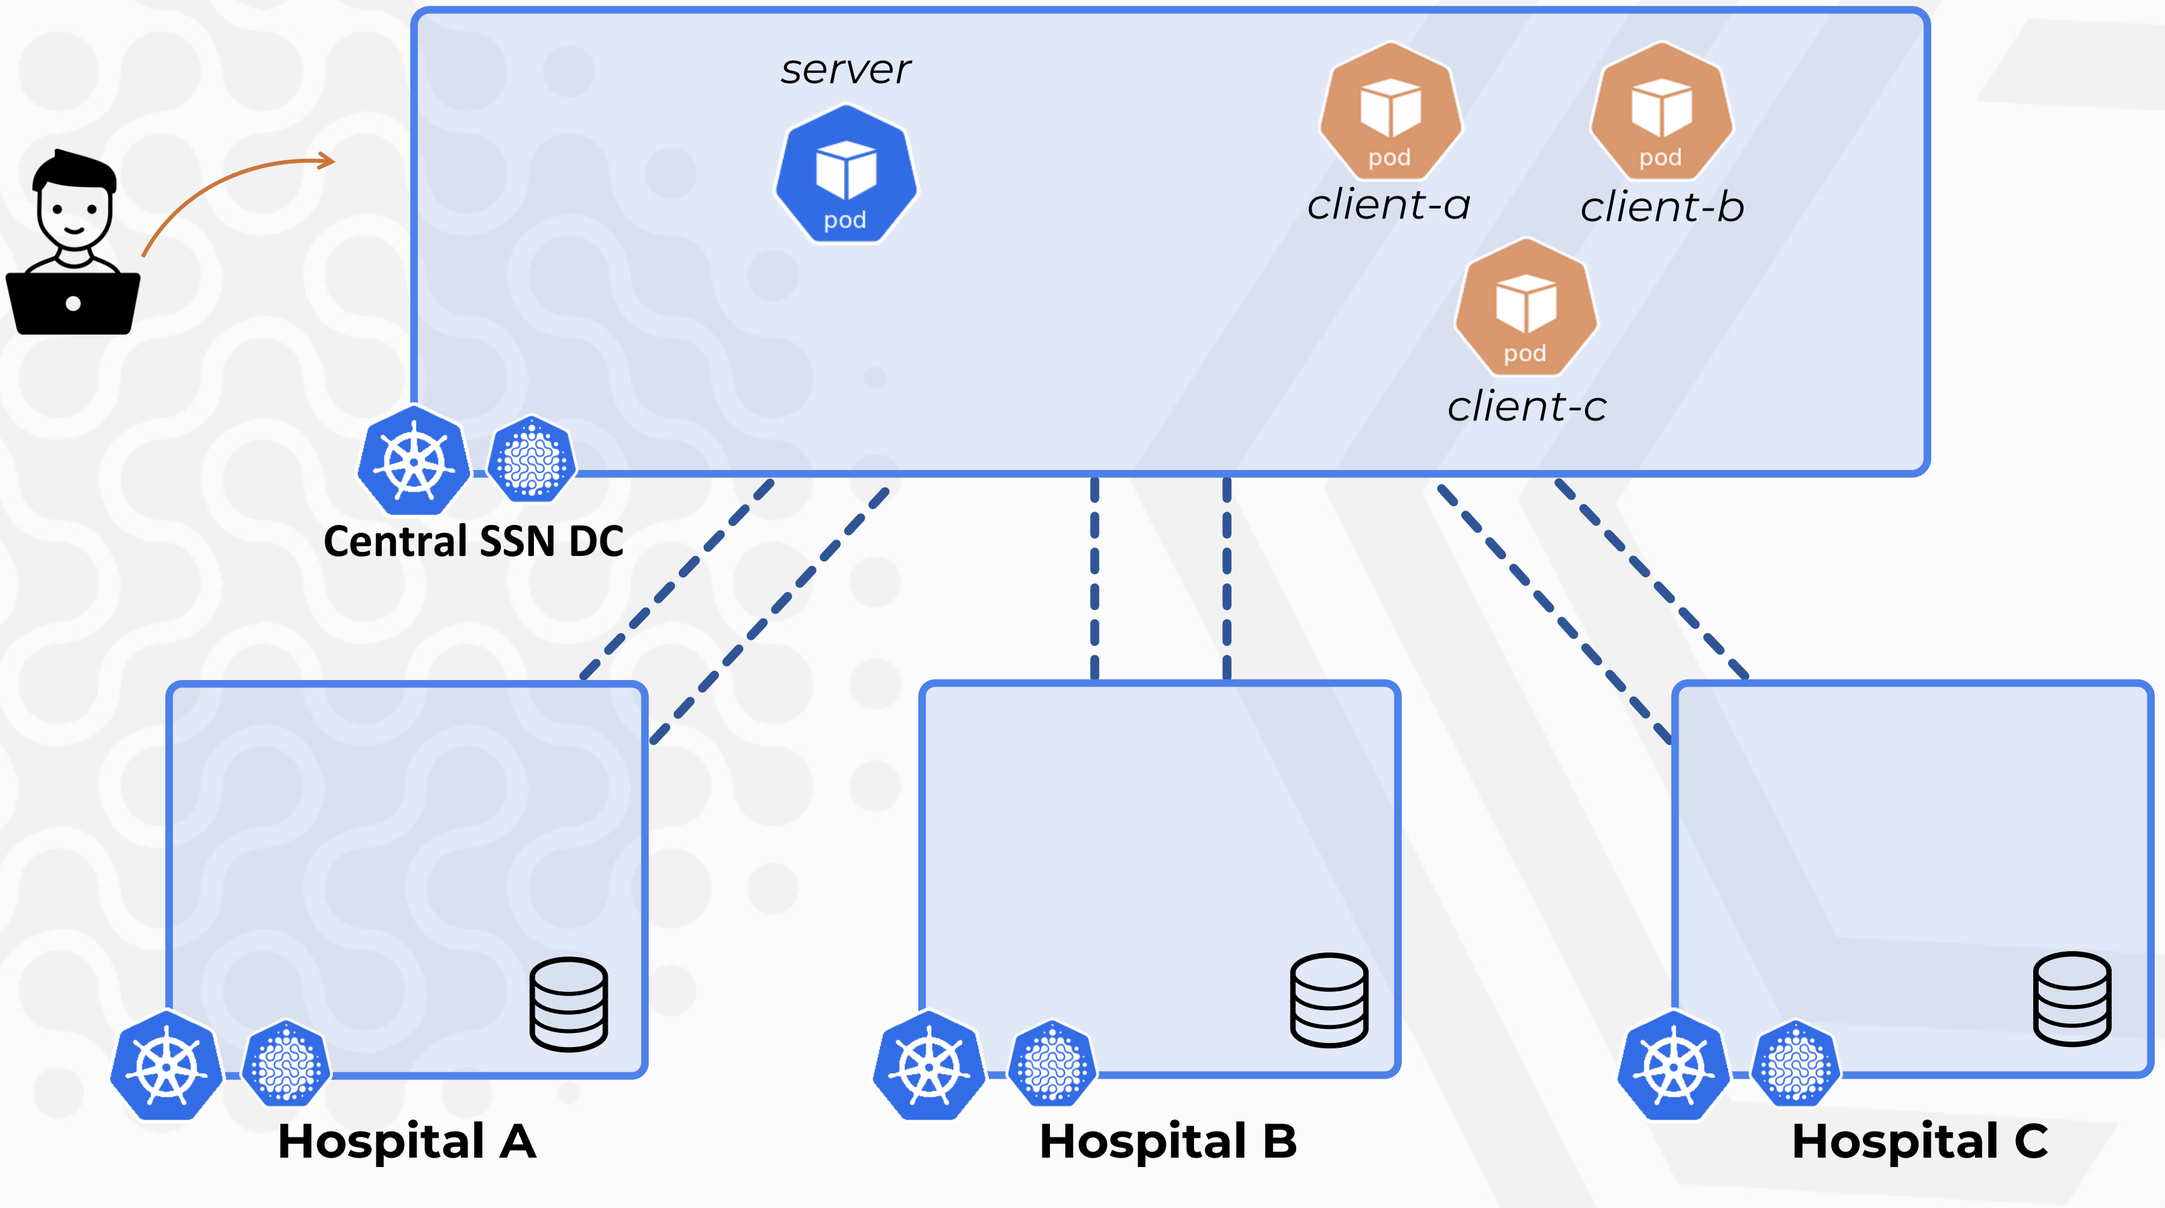

FL with Liqo and flower: demo

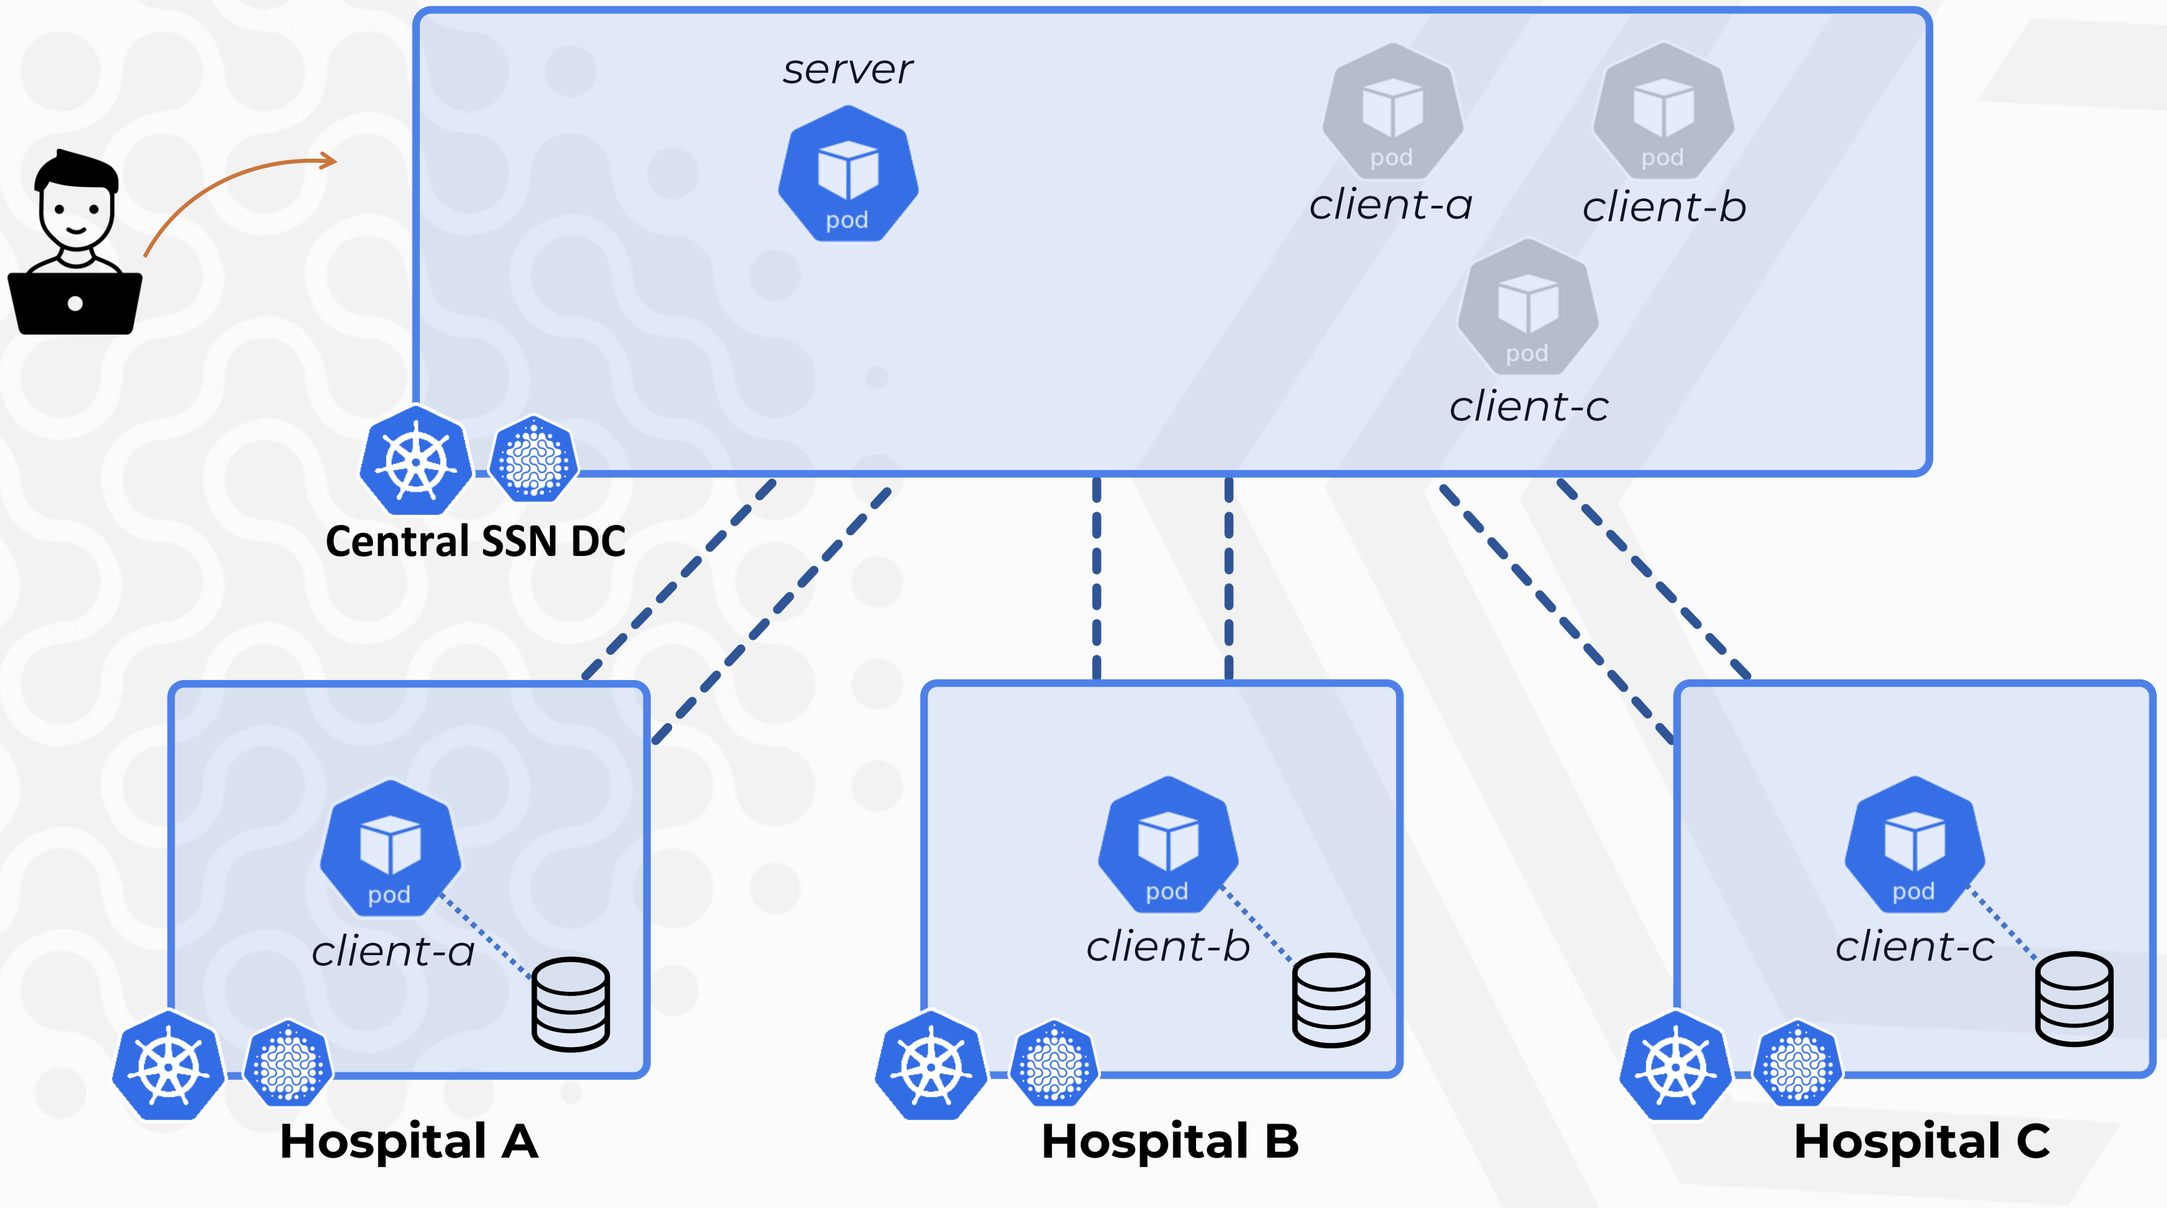

FL with Liqo and flower: demo

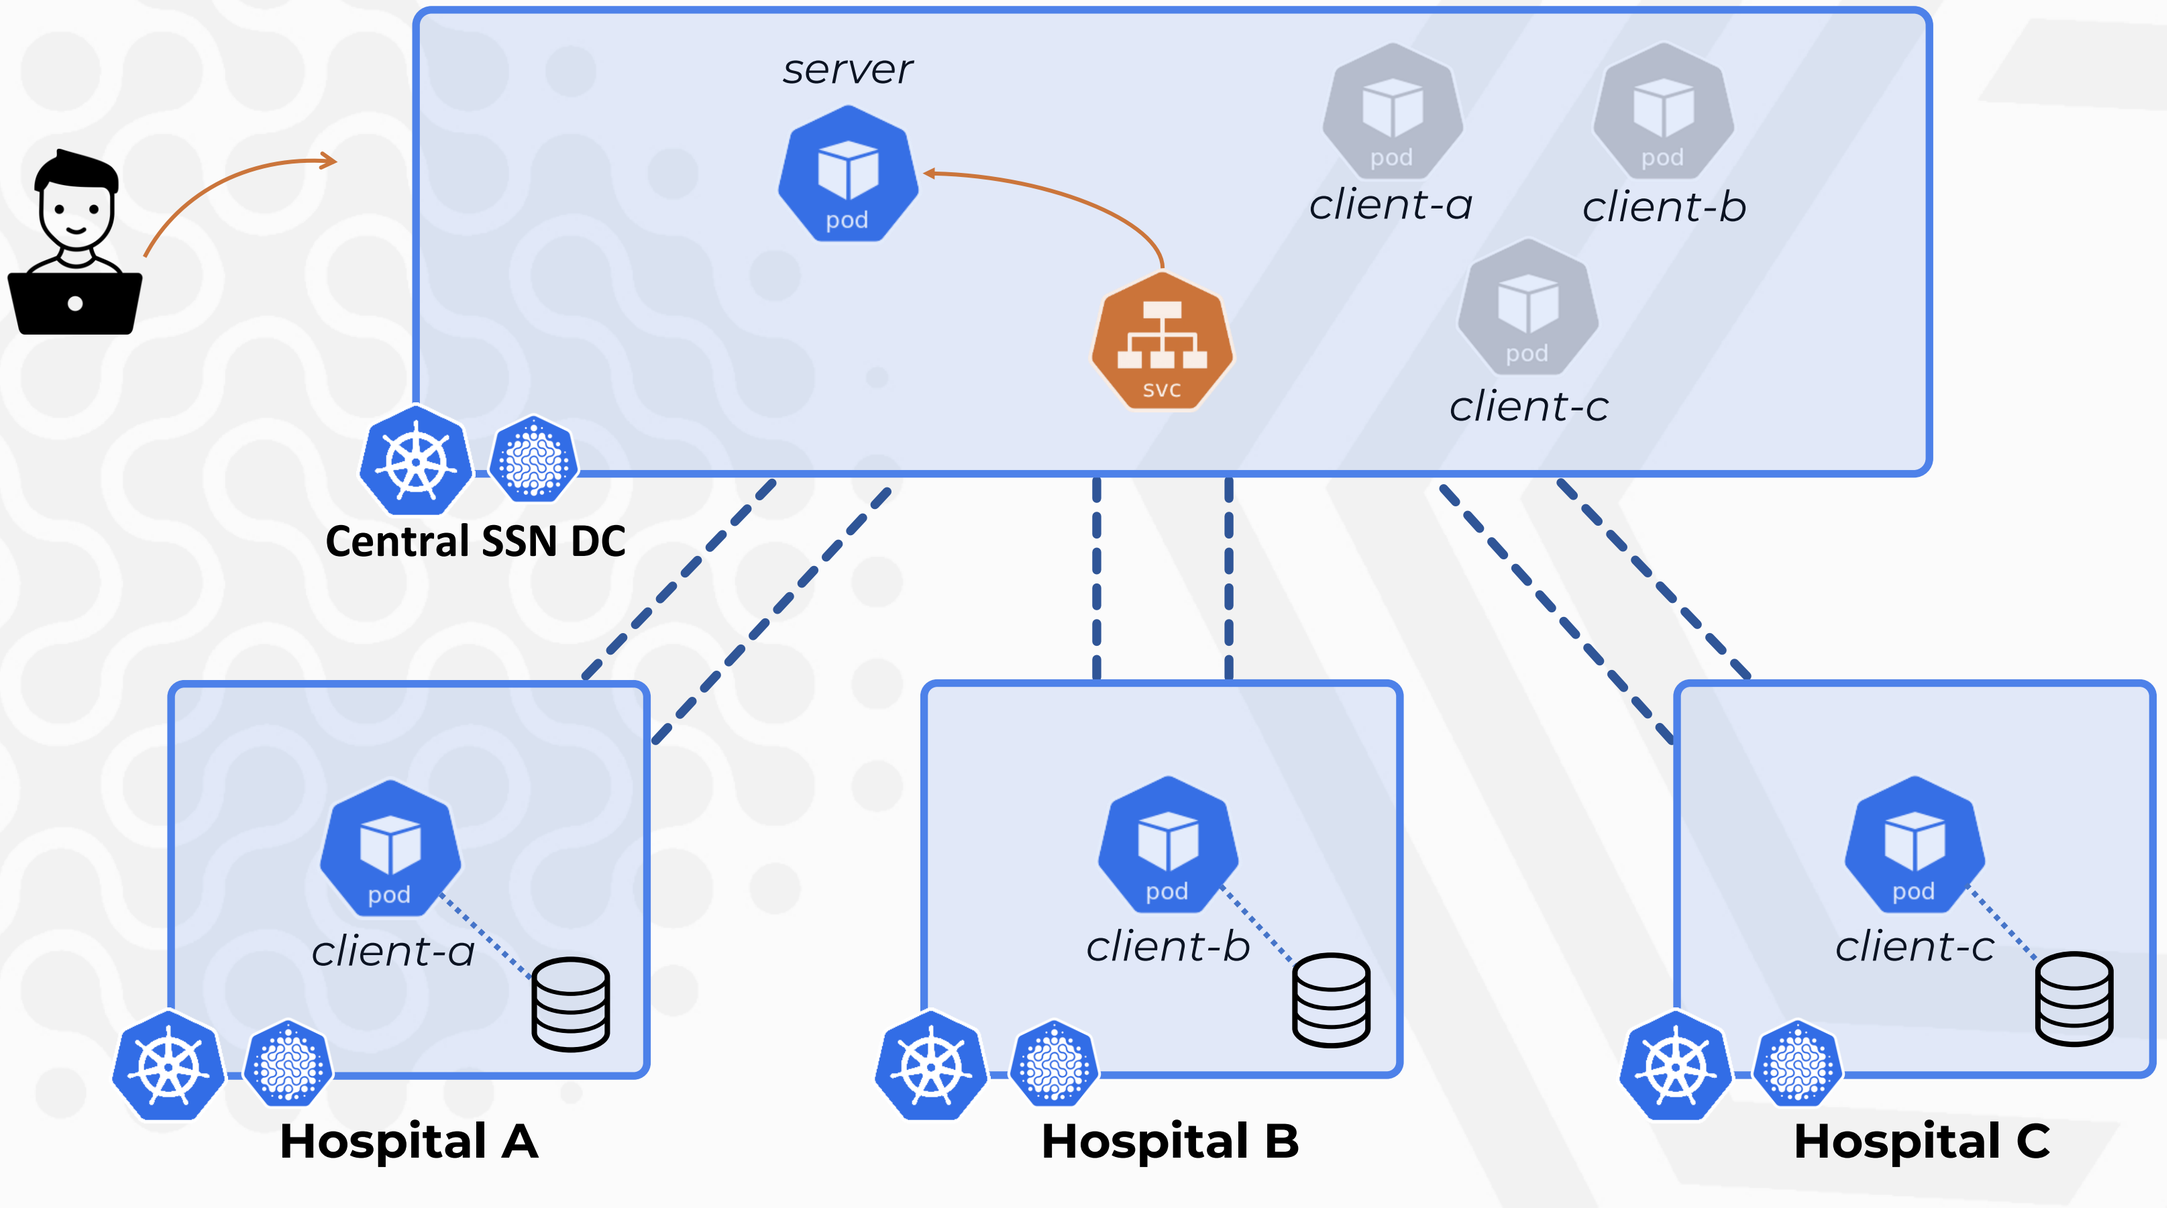

FL with Liqo and flower: demo

aruba.it KUBE

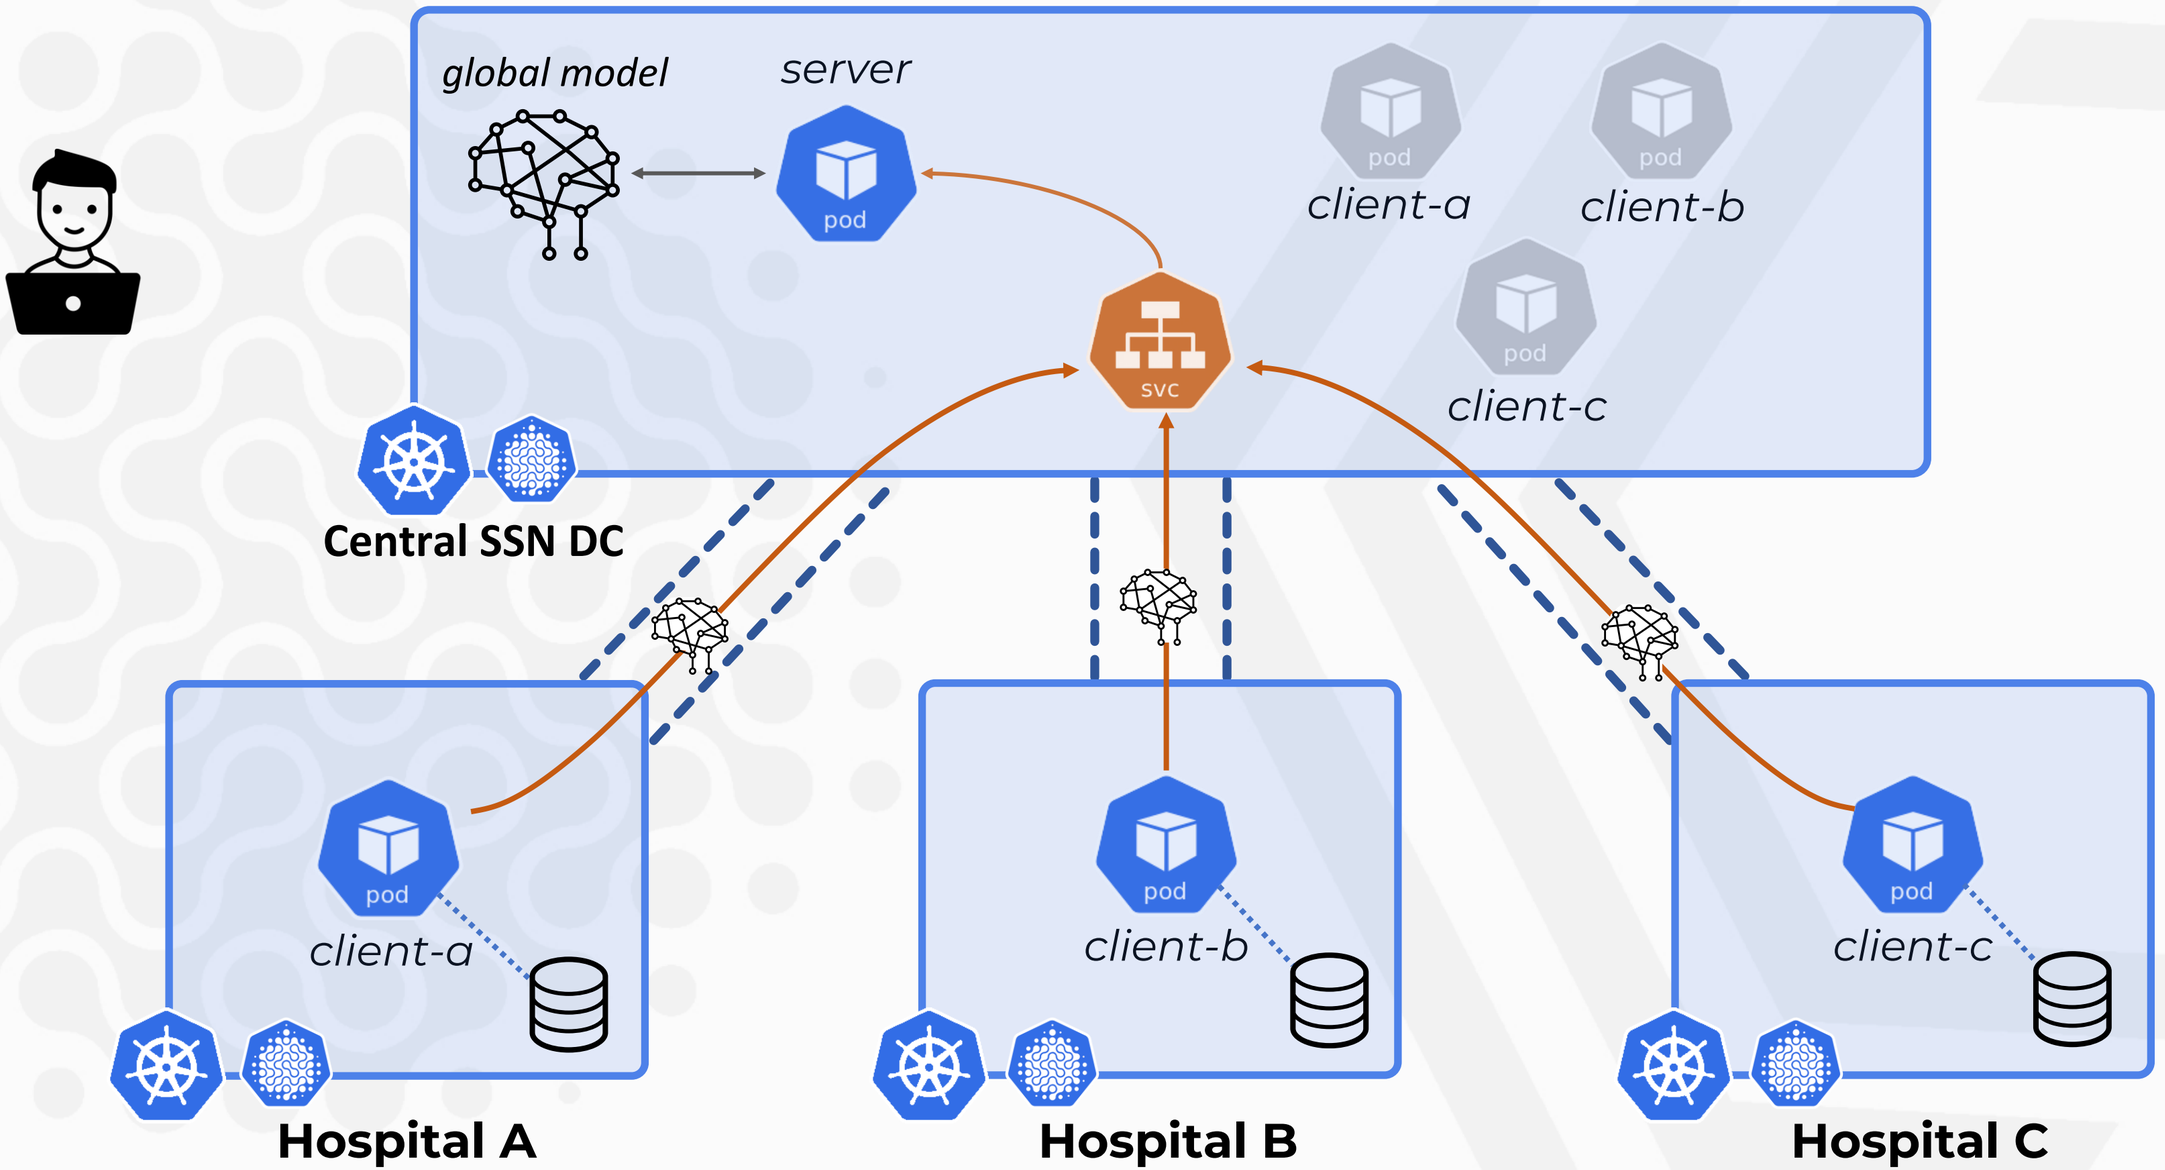

FL with Liqo and flower: demo

aruba.it KUBE

```yaml
apiVersion: v1
kind: Service
metadata:
  name: server
spec:
  ports:
  - port: 9092
    targetPort: endpoint-clients
    protocol: TCP
  selector:
    flwr.liqo.io: server
  type: ClusterIP
```

Server *(service)*

```yaml
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: client
spec:
  replicas: 3
  ...
  template:
    spec:
      containers:
        - name: node
          image: ghcr.io/...
          args:
            - --insecure
            - --superlink=server:9092
      runtimeClassName: liqo
```

Client *(statefulset)*

scheduled on Liqo virtual nodes

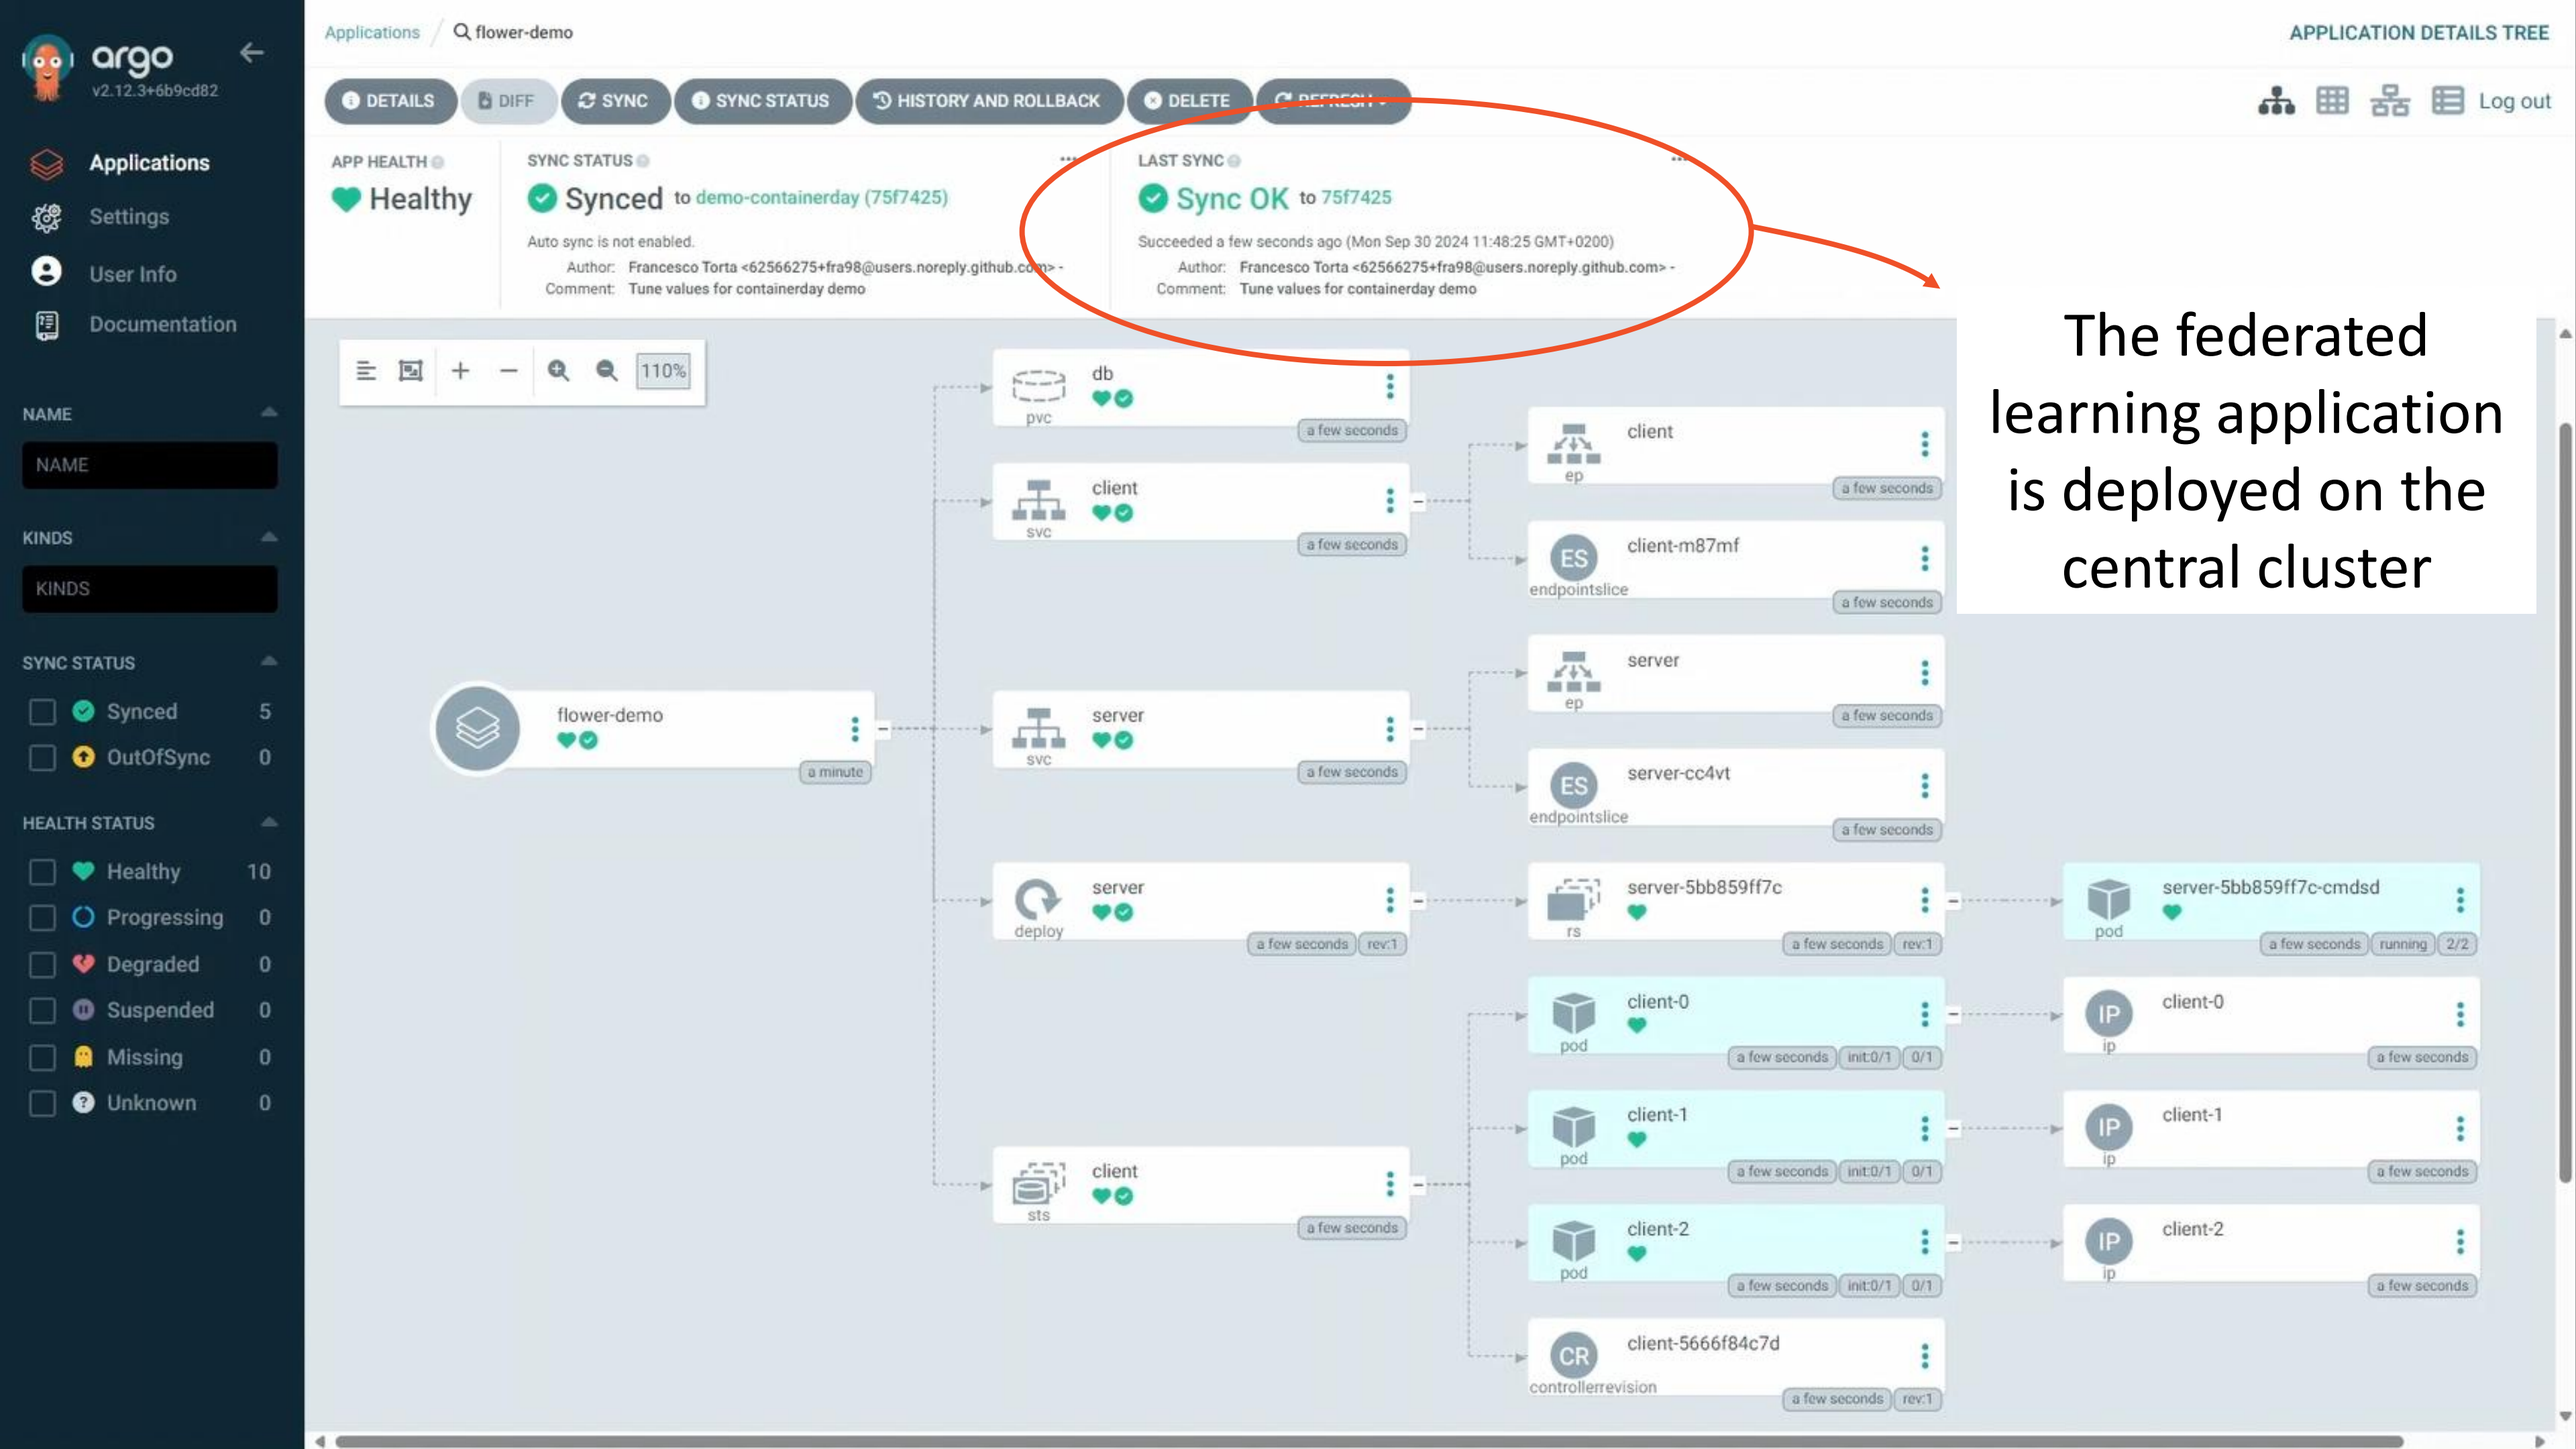

The federated learning application is deployed on the central cluster

FL with Liqo and flower: demo

Clients pods are scheduled on the virtual nodes (so they have been offloaded on the hospitals' clusters)

```
~ » kubectl get pods -n flower-demo -o wide
NAME                      READY   STATUS    RESTARTS   AGE     IP            NODE
client-0                  2/2     Running   0          7m58s   10.71.0.12    hospital01
client-1                  2/2     Running   0          7m58s   10.97.0.12    hospital02
client-2                  2/2     Running   0          7m58s   10.68.0.12    hospital03
server-d57469795-tfgl9    2/2     Running   0          18m     10.244.1.27   central-cluster-worker
```

Liqo Virtual Nodes

Conclusions

We obtained a solution coordinating different **entities, application and infrastructure**.

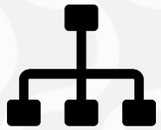

Flexible and **centralized** orchestration of workloads

The **infrastructure is transparent to applications**

Inter-cluster connection via **secure channels** out of the box

Built-in privacy: no data is moved, but **computation is moved to data**

Let’s stay in touch!

linkedin.com/in/claudio-lorina-0703b4237
linkedin.com/in/gzangari

Here you can find the GitHub repo with the demo and the references to the projects we used

https://linktr.ee/arkcloudday

Promosso da aruba CLOUD

Organizzato da {codemotion}

CLOUD BUDGET REVOLUTION: FINOPS UNLEASHED

Online Hackathon

Crea il futuro del budgeting intelligente per le risorse cloud

31 OTTOBRE - 28 NOVEMBRE | ONLINE

Crea il futuro del budgeting intelligente per le risorse cloud

31 OTTOBRE - 28 NOVEMBRE | ONLINE

Scansiona per il link all’iscrizione!

Thanks!

CLOUD DAY 2024

Milano, Nov 20

improove