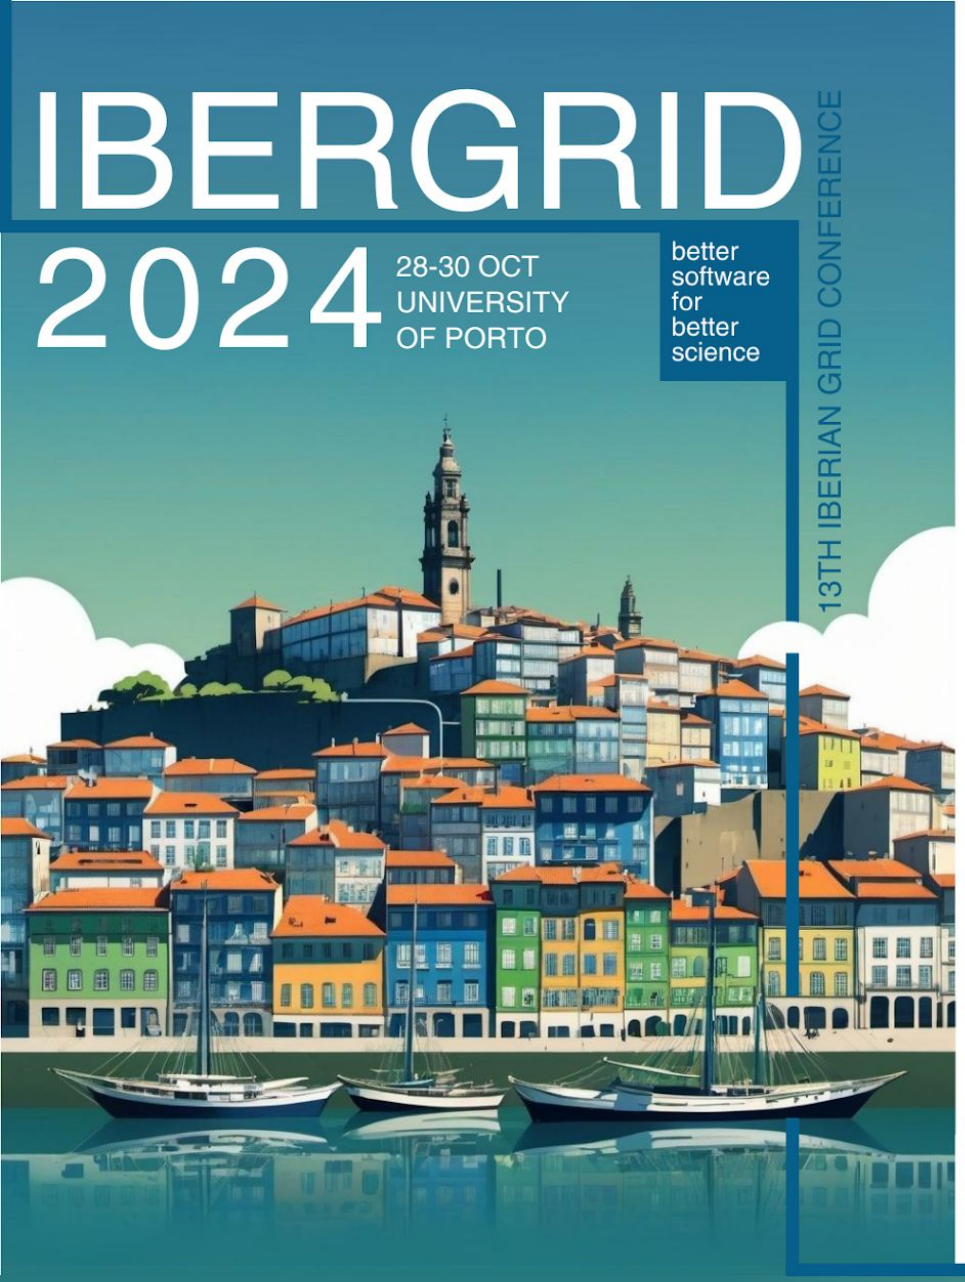

Cloud bursting to commercial providers with Kubernetes

Speaker: Samuel Bernardo
LIP and INCD

Overview

- Our Team
- Promotion and partnership motivation
- Improved sustainability with Google Cloud Platform (GCP)
- The challenge for software developers
- Kubernetes adoption to leverage expandability
- INCD integration with GCP

Our Kubernetes Team

Jorge Gomes

João Machado

Samuel Bernardo

João Martins

César Ferreira

João Pina

Zacarias Benta

Miguel Viana

Open Cloud for Research Environment

- Call for commercial cloud redistribution
- Extend to platform resources of commercial providers
- Distributing state-of-the-art digital services via European Open Science Cloud
- Promote research infrastructures moving forward to explore new mechanisms

- Service Aggregator
- Combined use of services
- Provide an easier solution for users willing to use commercial cloud for additional capacity or services
- Additional protection and support for the community

Commercial provider benefits

- Worldwide distributed resources
- Test brand new hardware (TPUs, GPUs, CPUs)
- Access to resources not available (expensive and less frequently used)
- Provider specific added value services

Dataflow
BigQuery
Cloud Storage
Vertex AI
Dataproc

Sustainable infrastructure partnership

Additional Capacity and Capabilities

Base Capacity and Capabilities

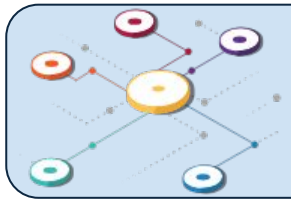

Research and academic community

Answer additional requests:
- hardware
- services
- extended limits

Safety and support:
- sustained usage
- users protection
- infrastructures abstraction

Computing resources:
- avoid lock-in
- minimise costs
- burst into Google Cloud

Architecture towards sustainability

- Container-based approach migration of existing software
- Event driven solutions
  - Function as a Service
  - Task execution
- Serverless adoption
- Minimise costs
- Users support over the software migration

WORSICA OPENCoastS OpenEO SQAaaS

Kubernetes adoption by developers

- Start with docker image creation
- Prepare the deployment using docker compose
- Use available tools to help developers to generate kubernetes ready deployment configurations

Kompose - docker compose to kubernetes

- configuration files
- helm chart

SQAaaS - Quality assessment and deployment automation

Kubernetes adoption by developers

- Use available tools to help developers to generate kubernetes ready deployment configurations

ArgoCD: GitOPS implementation over **Kompose** generated configurations (CD)

Kubevela +  **FluxCD**: Kubevela GitOPS application (CD)

Scaffold: build, push, test, deploy, verify (CI/CD)

Google Kubernetes Engine (GKE)

- Autopilot
  - most favorable pricing model (tailed to the application)
  - kubernetes core components are managed by Google
  - limited access to the kubernetes nodes (shared)
- Standard
  - pricing model adds the nodes allocation to the application operational costs
  - dedicated infrastructure for the nodes
  - support custom node configurations

GKE Autopilot on-demand pricing model

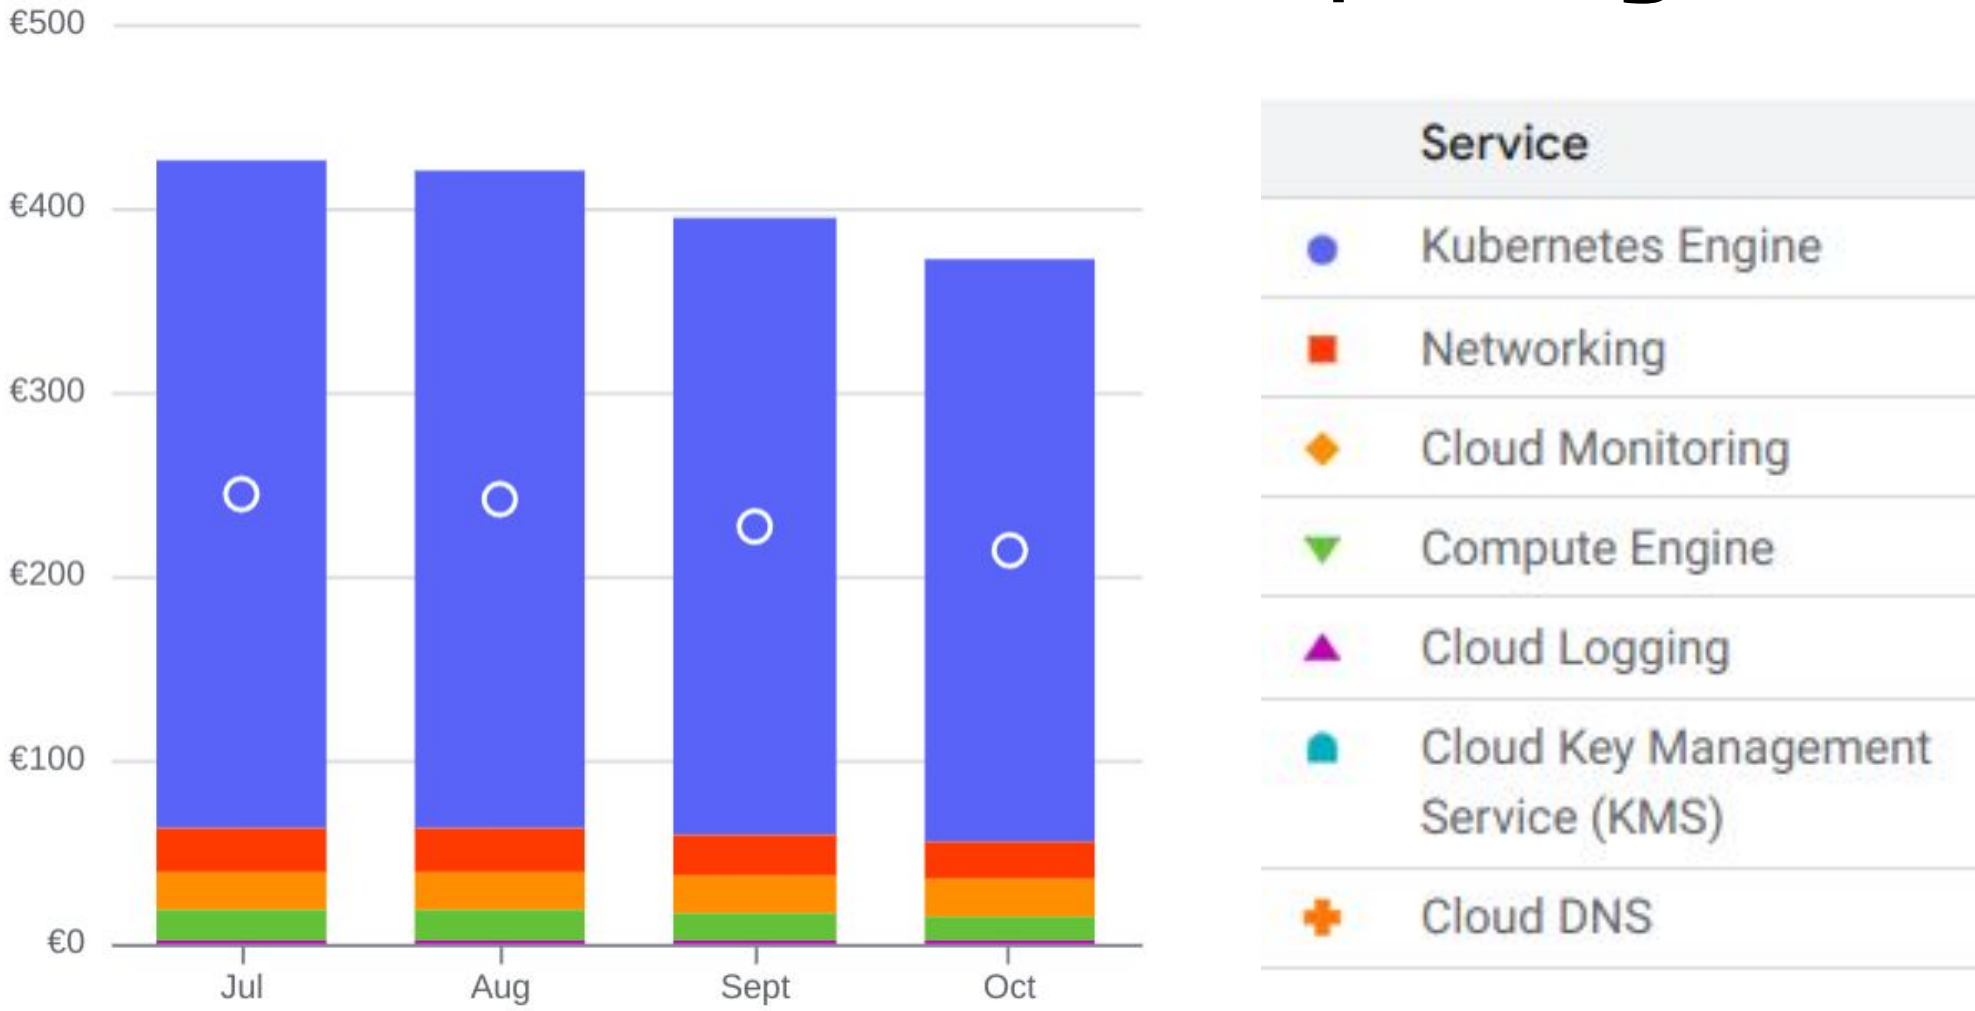

GKE Standard on-demand pricing model

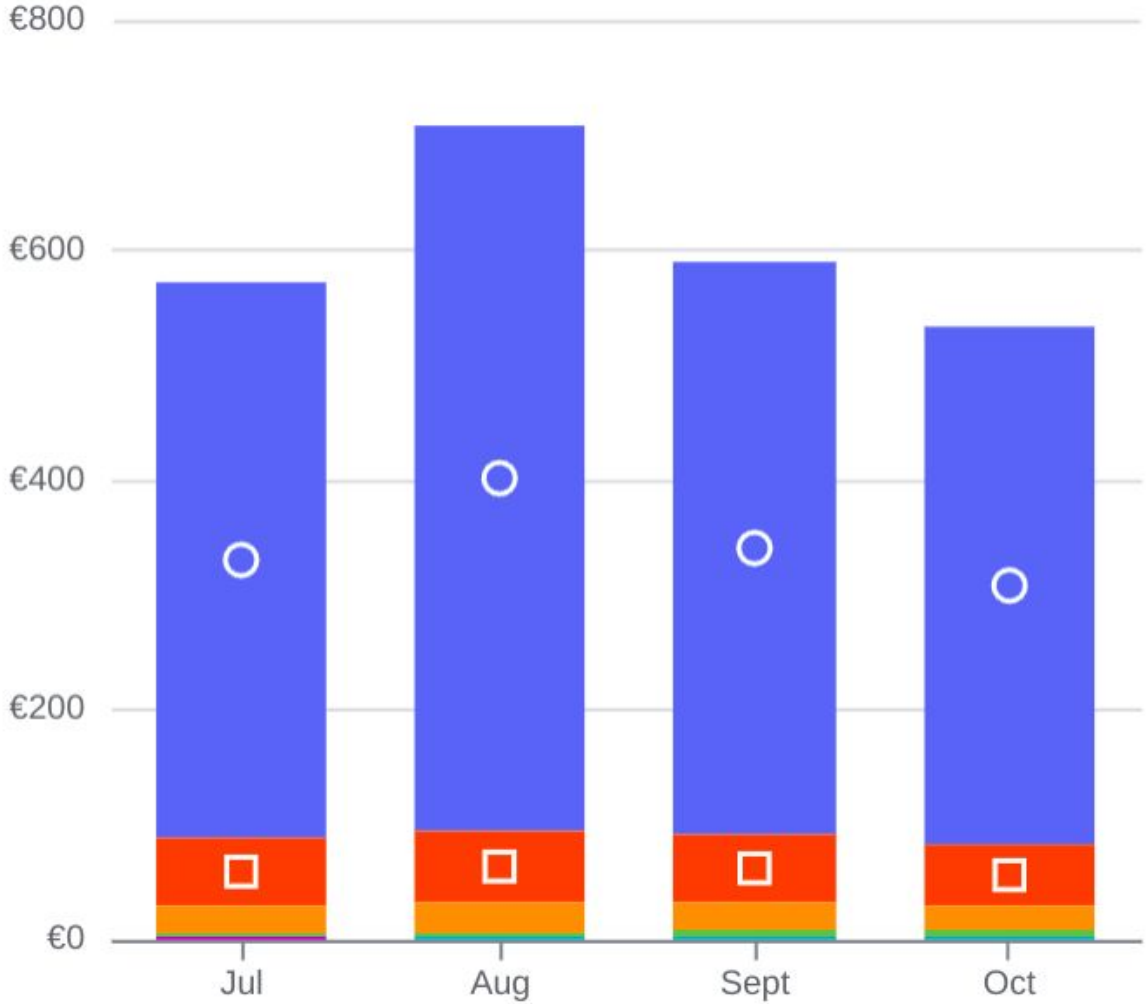

Service
Compute Engine
Kubernetes Engine
Networking
Cloud Storage
Cloud Key Management Service (KMS)
Cloud DNS
Cloud Logging
Cloud Monitoring

GKE projects organization

- Multiple different projects supported in same organization
- Projects can be placed in folders
- Folders allow to share roles with subfolders and projects

Autopilot GKE

Standard GKE

GKE service routing

- Ingress Controllers implementations (proxy): Nginx, Traefik
- Ingress Controllers can be shared or dedicated to each service

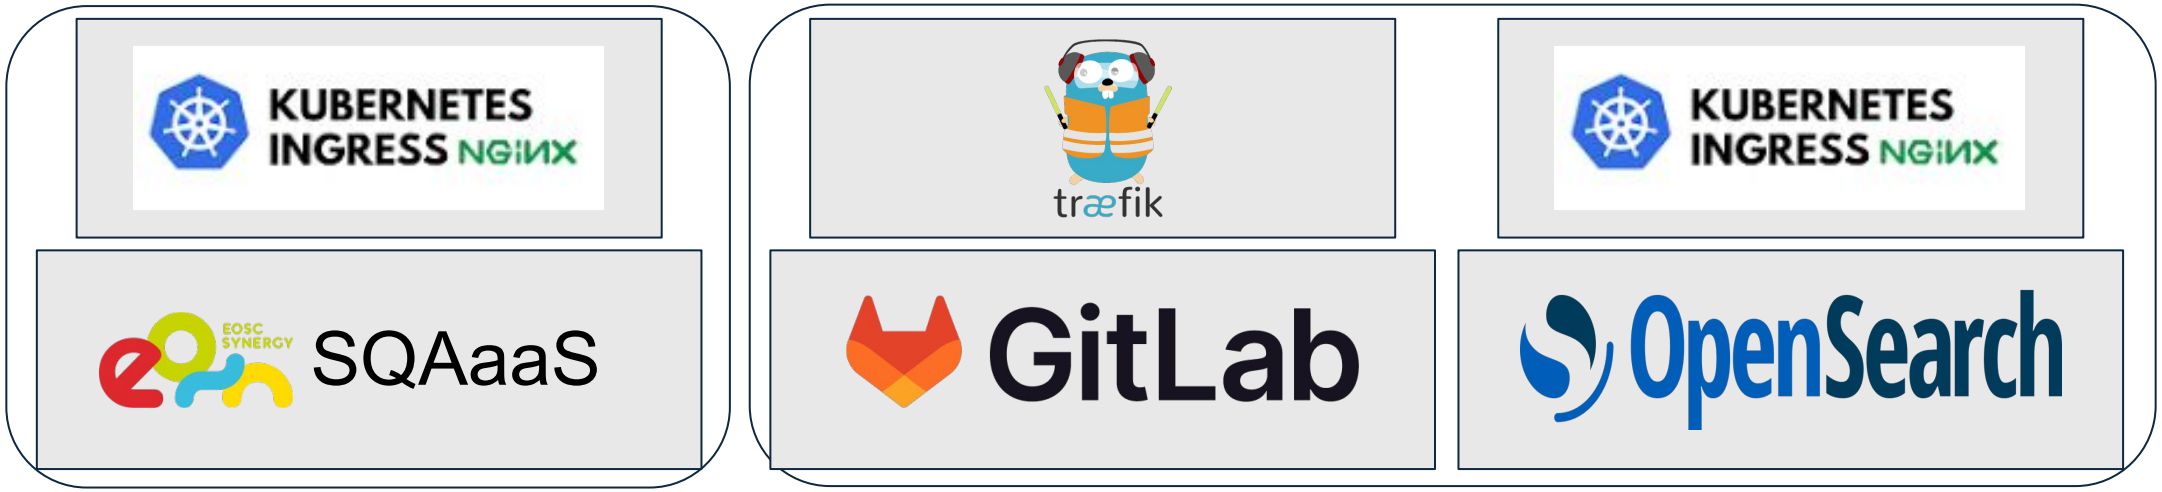

GKE service routing

- Each ingress controller have its own gateway (load balancer)
- GCP Gateways only allow traffic managed by INCD

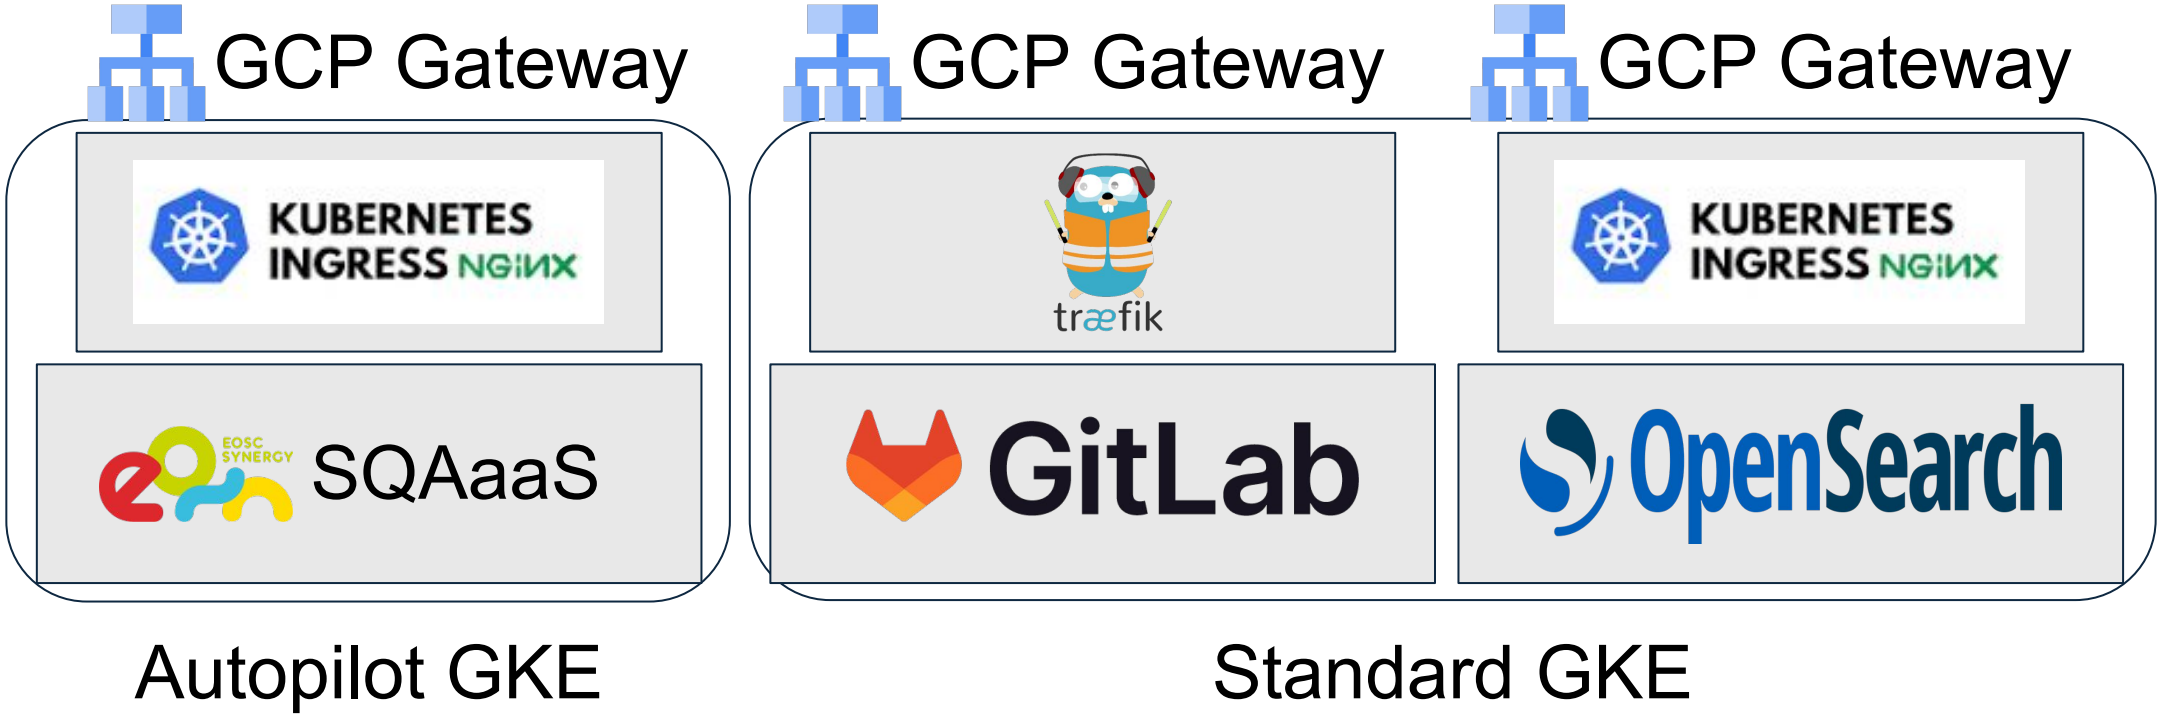

Using Google Cloud DNS for GKE

Cloud DNS - Google’s worldwide network distributed DNS

- INCD organization DNS zone (a delegation of incd.pt)
- Group (folder) or project DNS zone (as delegations of above)
- Use wildcard CNAME to forward requests to the Gateway
- These name records are only used for INCD routing to GCP

INCD delegated DNS zone for GCP

- Delegates routing domains to Cloud DNS
- Service DNS names for the users projects are only at INCD DNS servers
- Possible integration with FedCloud Dynamic DNS (IISAS)

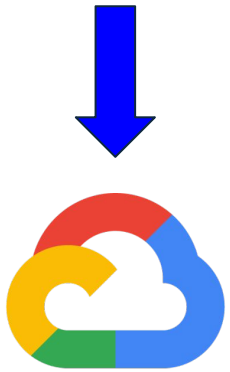

Google Cloud

INCD
org zone

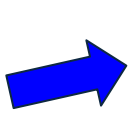

Group zone ■ ■ ■

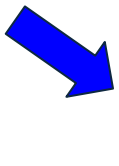

Project zone

INCD gateway for Kubernetes clusters

- On-premises ingress over Kubernetes dedicated infrastructure
- Use external names dynamically managed at GCP
- GitOPS methodology compliant
- Easy to manage with same configurations (templating)
- Allow data to be kept at INCD (important for customer protection)
- Data movement mostly from INCD site to external provider
- Minimise future data storage and transfer costs
- Services can move between providers with almost zero downtime
- Network traffic control through INCD infrastructure

INCD integration with GCP

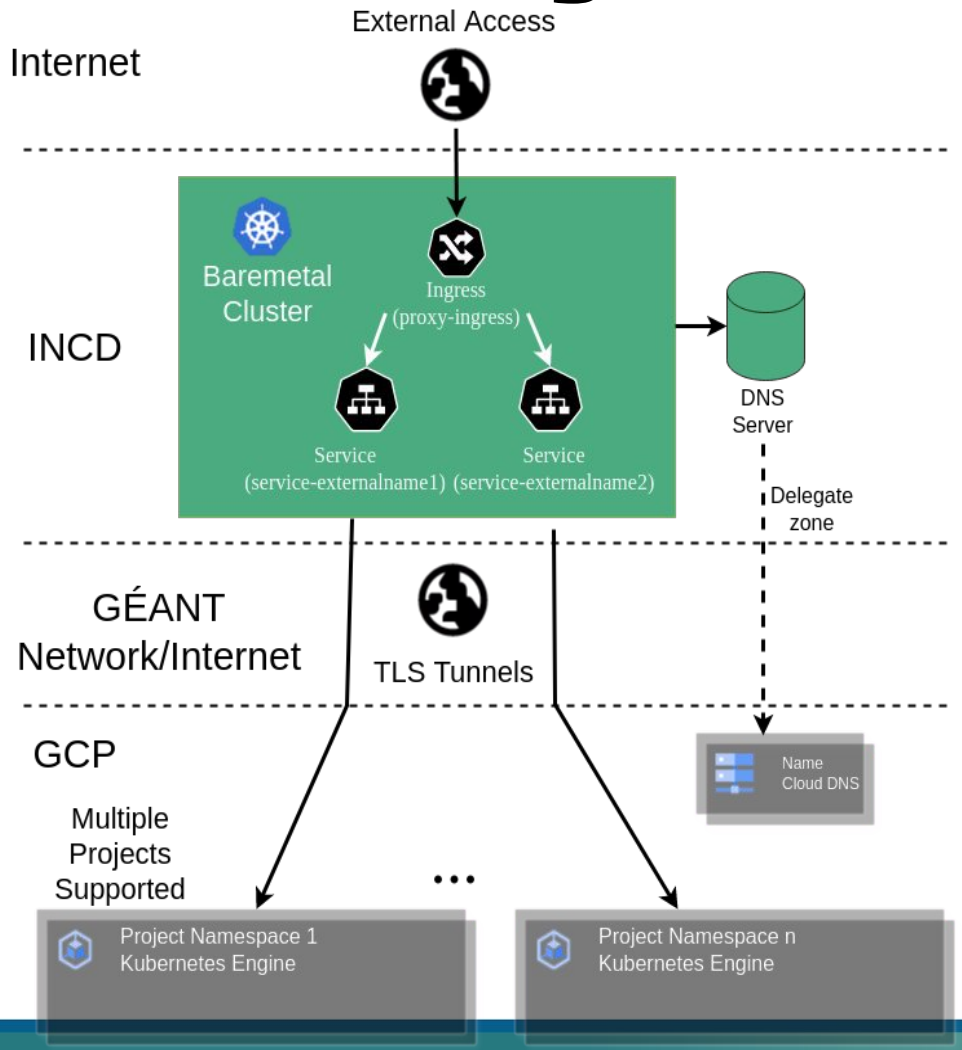

- Endpoints are exposed by INCD
- Endpoints are managed by proxies
- DNS via zones that can be delegated
- Projects can be moved to GKE
- Depending on project access can be:
  - Directed to INCD Kubernetes
  - Directed to Google GKE
- Tunnels between INCD and Google
- Multiple different projects in GKE

Thank you

https://www.incd.pt

helpdesk@incd.pt