
MariaDB
OPEN WORKS
BE UNSTOPPABLE

BETTER TOGETHER: MARIADB SKYSQL RUNNING ON GOOGLE CLOUD PLATFORM

AKSHAY RAM, SR PRODUCT MANAGER, GOOGLE CLOUD PLATFORM
NAMAN SHAH, PRODUCT DIRECTOR, MARIADB

Data listed as top growth area in CNCF Survey 2022

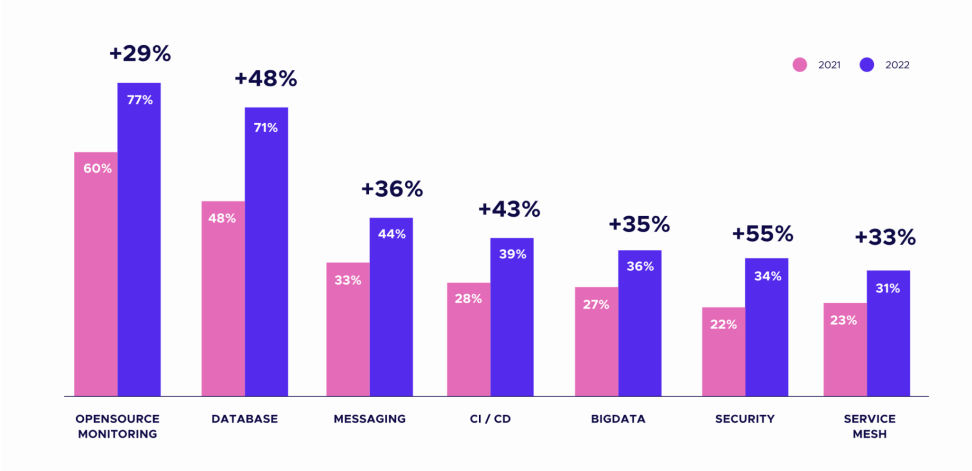

- **Databases** has the highest growth at 48%
- **Messaging and Big data** also data apps with Kafka, Spark…

Source: CNCF 2022 Survey

STATEFUL APPLICATIONS ON GCE

Stateful applications on Kubernetes

According to the Data on Kubernetes report, 70% of Kubernetes customers run stateful Apps.

Top Stateful Applications	
MariaDB	RabbitMQ
Redis	Postgres
MySql	Kafka
Mongo	Elasticsearch
Elastic	Cassandra

Exponential growth as customers see benefits on GKE

Exponential growth with stateful clusters doubling each year on GKE

Google now a gold sponsor on Data on Kubernetes (source)

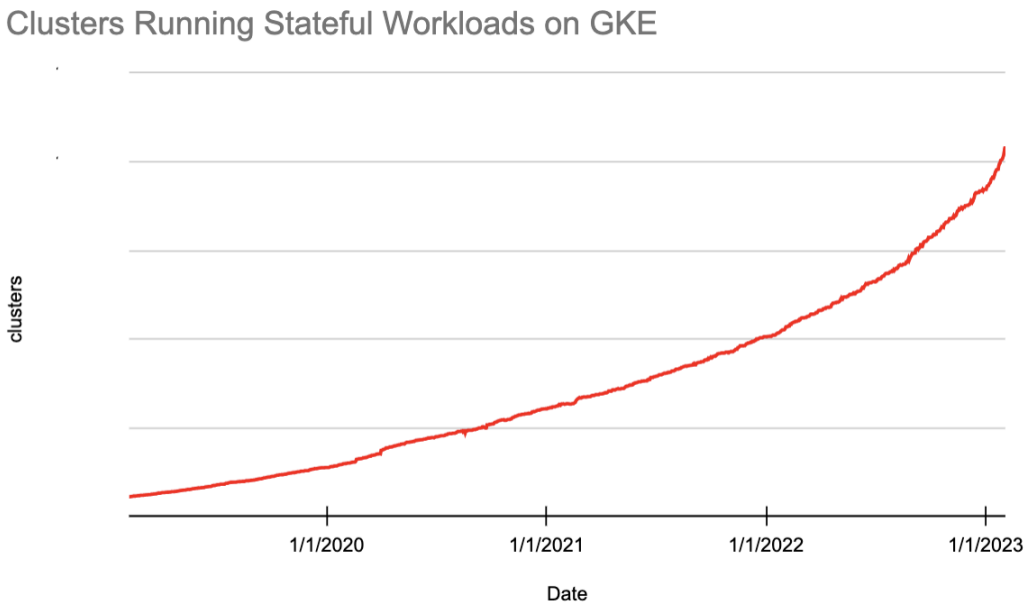

Spectrum of Stateful Apps on GKE

Do it yourself (DIY)

Eg. MariaDB, postgresql

Apps deployed as container images and managed by customers

Kubernetes Operator

Eg. Elastic operator

Apps deployed as container images with management shared with operator contracts.

Data SaaS

Eg. MariaDB SkySQL

Apps that are fully managed Saas solutions for end users

Self Managed	Partially Managed	Fully Managed

Why DIY customers choose GKE

Built in orchestration	GKE for stateful apps gives customers out of the box CI/CD, config as code and stateful primitives, built-in backup
Organizational gravity	Running both stateless and stateful apps on Kubernetes removes silos and leverages economies of scale in organizational expertise
Ecosystem	Operators allow simple one click install a database on our GKE clusters.

Why SaaS customers choose GKE

Multi-cloud infrastructure	The need to run anywhere, including across clouds, is a salient feature of a SaaS
Built-in capabilities	Out of the box cost allocation for unit economics, built-in capabilities like Backup for GKE, blue-green for safe deployments
Empathy	Saas uses use Kubernetes themselves for their stateless applications and running on Kubernetes builds empathy

GKE Stateful Platform

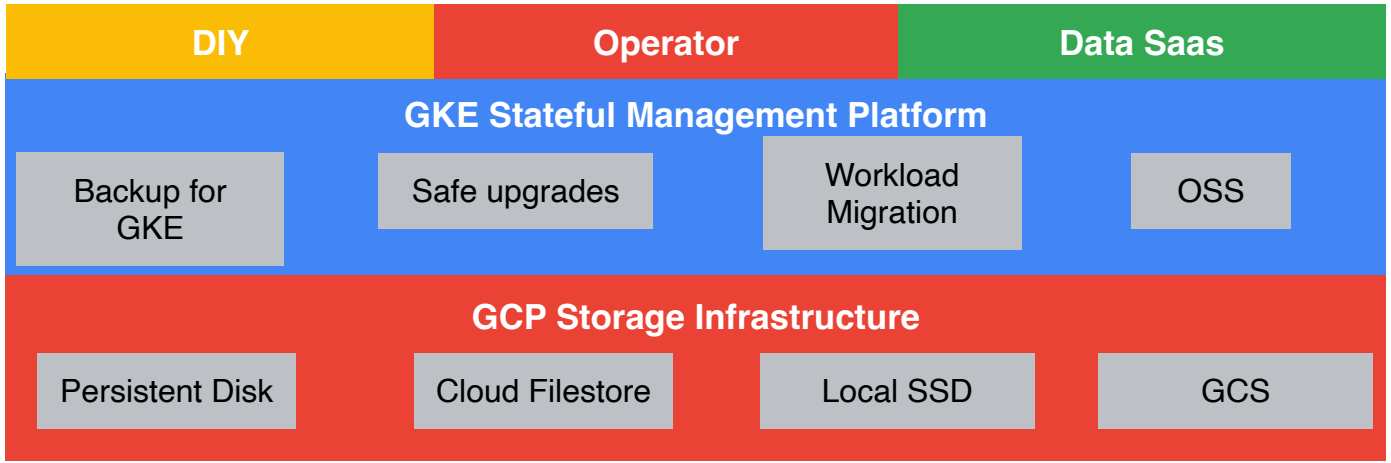

Built-in automated backups

Backup for GKE

Fully managed backup and restore plans

- Fine grained controls for automated backup plans by cluster or namespace
- IAM based policy control
- Restore within clusters or across clusters

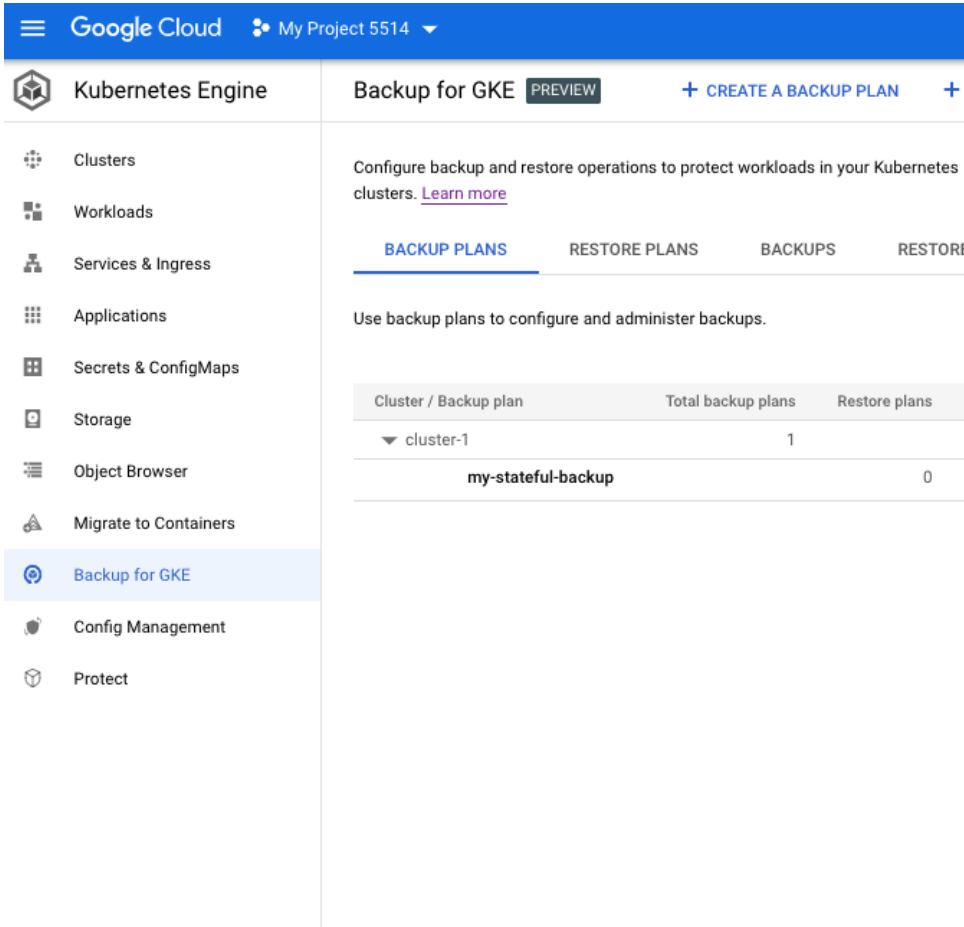

Upgrade stateful apps with peace of mind

Safe upgrades

Extensive controls for safe rollout

- Maintenance and exclusion windows
- Built-in blue-green deployments with rollback
- Kubernetes pod disruption budgets

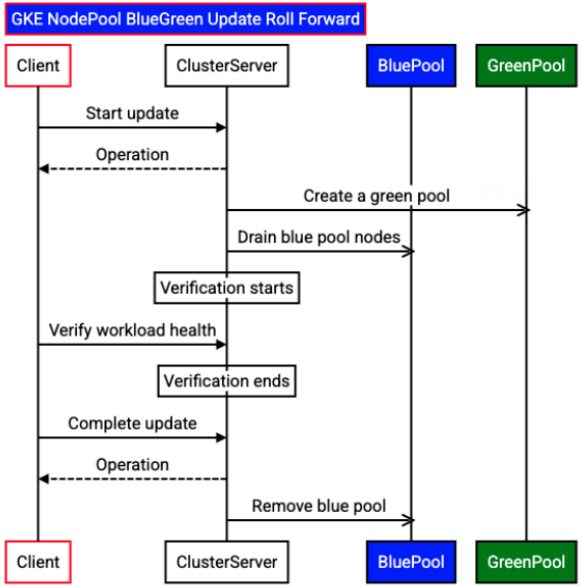

Extensive selection w/ GCP storage integrations

	 **Cloud Filestore**	 **Persistent Disk**	 **Local SSD**
Storage Description	Persistent storage for applications that need NFS and multiple readers and writers	Persistent storage for good price-performance and flexible sizes	Performant Ephemeral storage
Capabilities	Regional by default with Filestore Enterprise.	Zonal with option of read only regional	Zonal
Applications	Content management systems, rendering and media processing	Databases, Boot disk, Message queues, Cloud IDE	AI/ML, Analytics, Caches

Takeaways

01

Day 2 ops is a day 1 consideration

Set up pod disruption, safe rollout best practices along with observability to ensure streamlined day 2 ops

02

Storage choice based on use case

Choose a storage class that best suits your application needs eg. PD for databases, Filestore for content management, Local SSD for AI/ML

03

DIY or SaaS based on level of control

You can DIY for more control over the stateful app or just use a SaaS and have everything managed

04

GKE for modern stateful applications

Use out of the box capabilities to ensure best practices for stateful apps to keep them up to date, available, observable and protected. Use GKE Autopilot for added simplicity

SKYSQL ON GOOGLE CLOUD

MariaDB

SKYSQL ON GOOGLE CLOUD

- 17 global regions in SkySQL
- Leveraging wide range of services
- Certificate Authority service
- SkySQL is built on Kubernetes and leverages unparalleled GKE capabilities

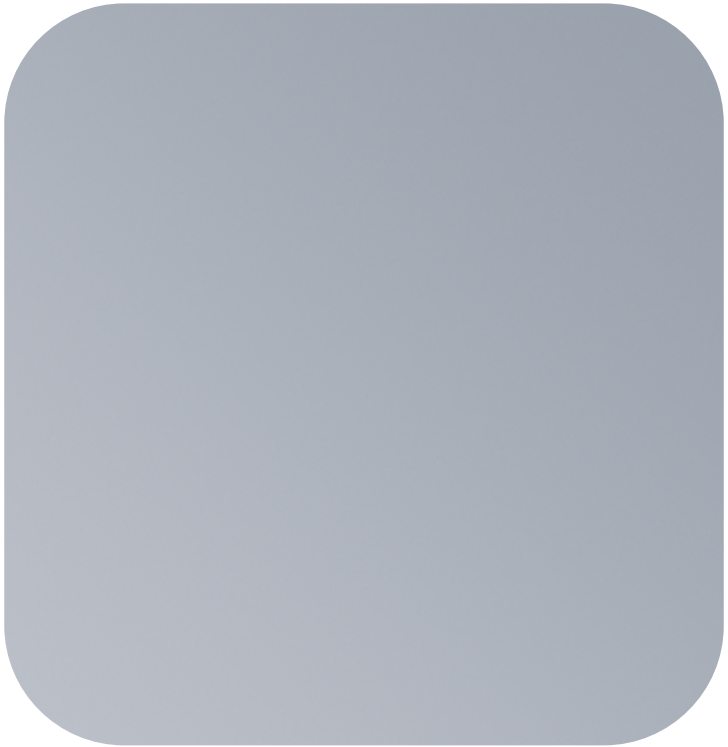

SKYSQL ON GKE

Build once, run anywhere

Kubernetes gives us the ability to support multi-cloud deployments and enables us to cater to our customers across clouds

Modern Orchestration

Running on a capable platform with built-in features for stateful applications allows SkySQL to scale, deploy safely and manage

Customer Empathy

SkySQL customers run their stateless applications on Kubernetes. Running SkySQL on GKE helps us build the same vocabulary and architecture familiarity with our customers

SKYSQL ARCHITECTURE

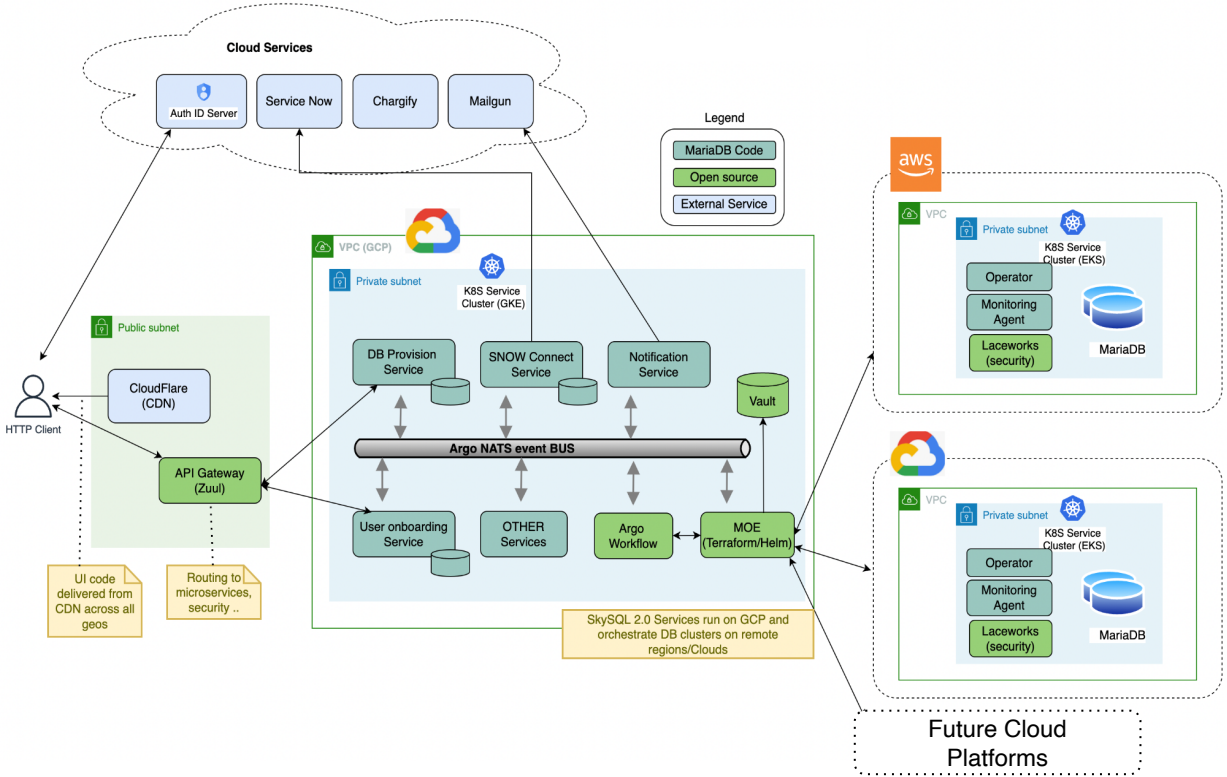

SkySQL Cloud Database Service Architecture

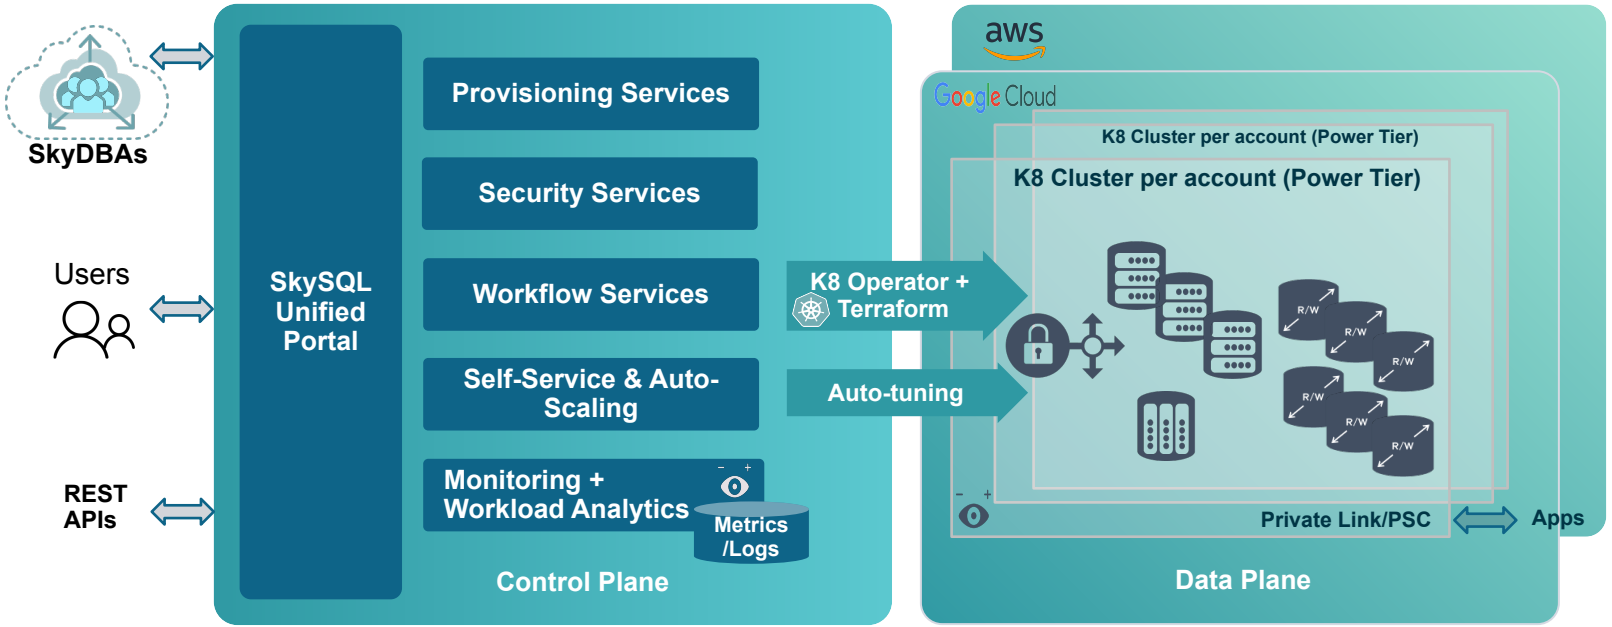

Legend

- Kubernetes Cluster (AWS – EKS, GCP – GKE)
- MaxScale Intelligent Data Proxy
- Enterprise Server
- ColumnStore Engine
- Xpand (Distributed SQL)
- Observability services

SkySQL Cloud Database Service Architecture (Detailed Version)

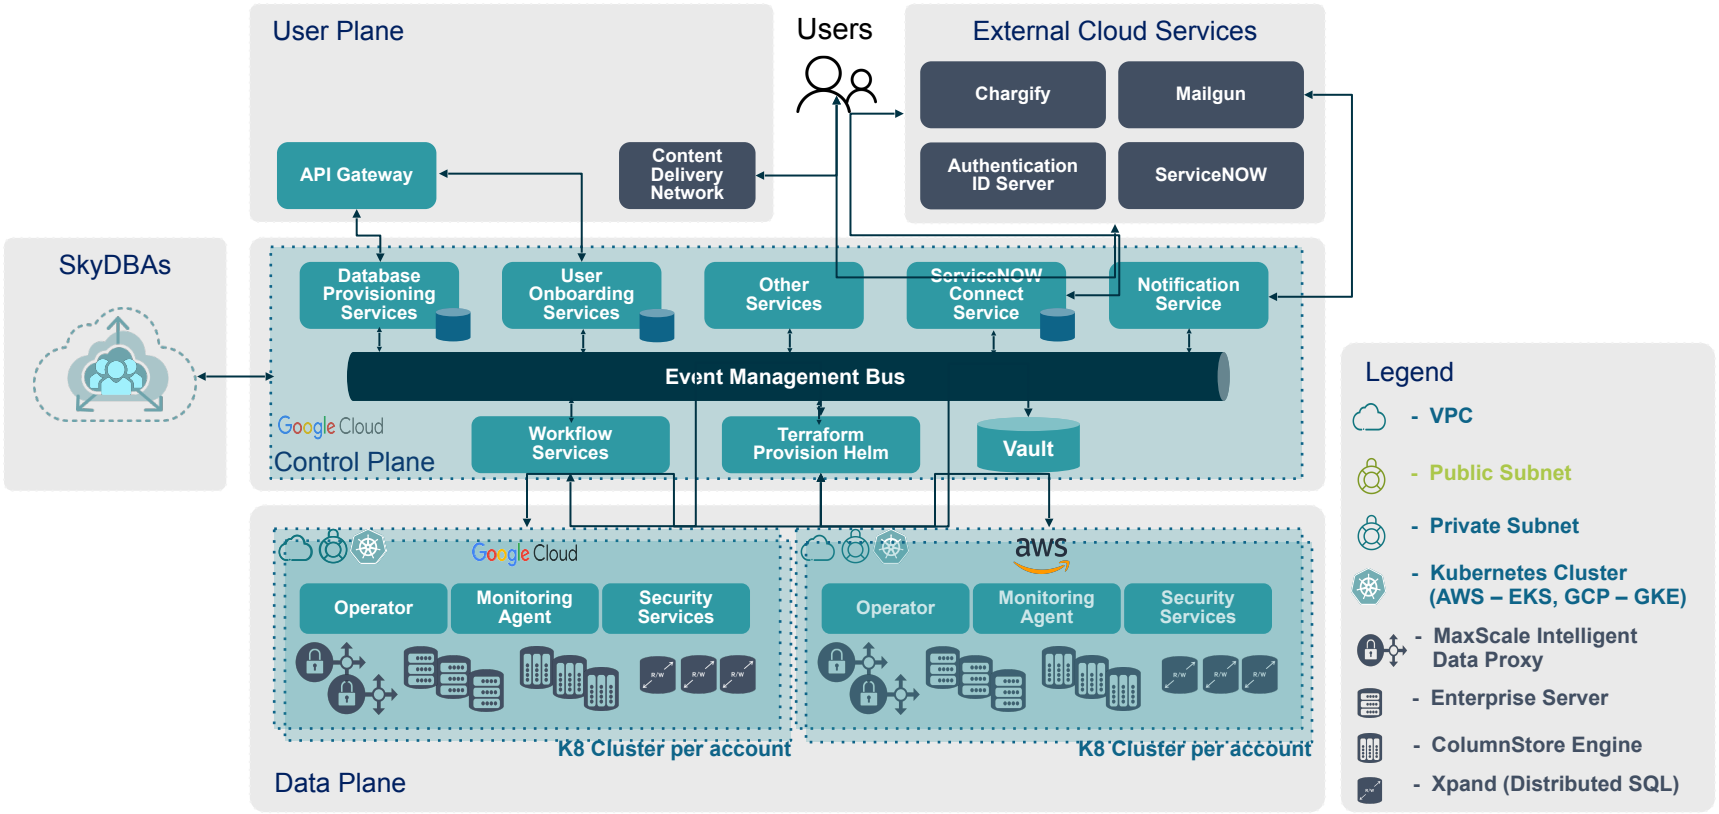

MariaDB

Operational Ease

- Ease of managing hundreds of clusters and their control planes
- GKE UI and API to manage our global workloads, services, ingress and config maps
- Kubernetes control plane and nodes upgrade without impacting availability

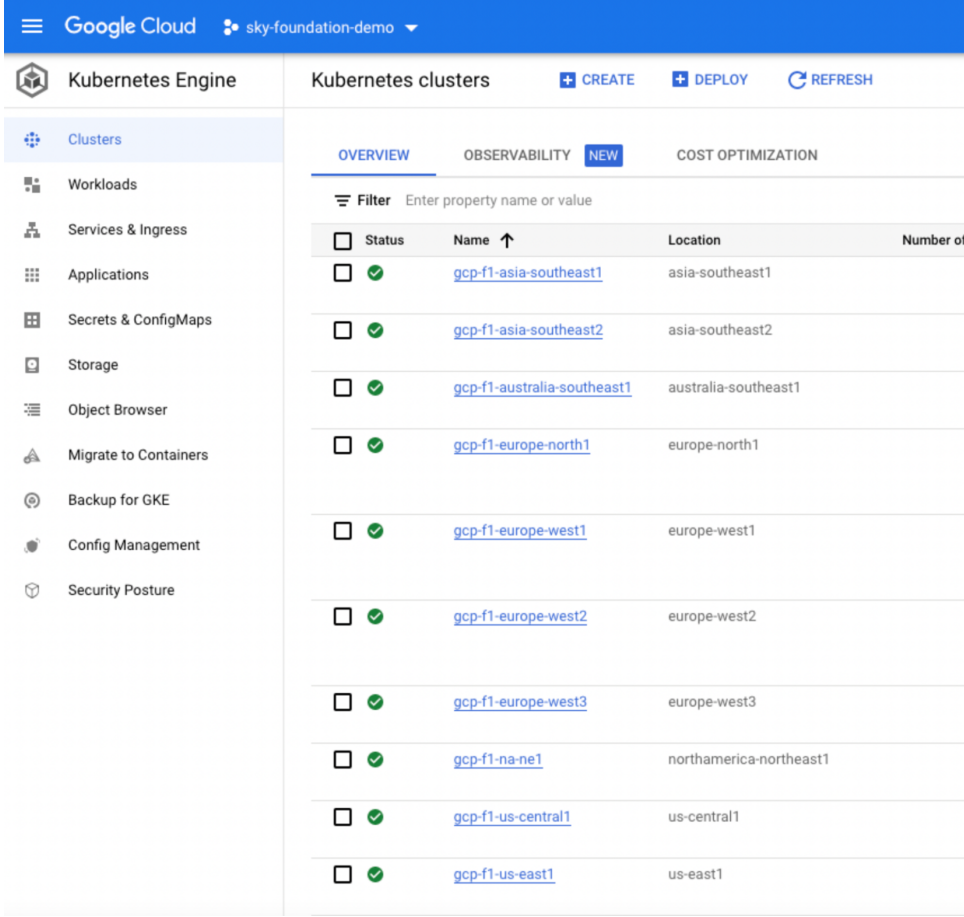

NEXT STEPS

Check out these sources to learn more about MariaDB

- **OpenWorks sessions to watch**
  - SkySQL vs. AWS RDS vs. GCP Cloud SQL
  - Fireside chat: Best Practices for Migrating Your On-Premises MariaDB Deployment to SkySQL
- **Visit the Marketplace today**
- **Create your Google account today**
- **Try SkySQL for free**
  - Try the full SkySQL service with a $500 credit, including ticketed support
  - https://mariadb.com/products/skysql/get-started/


MariaDB
OPEN WORKS
THANK YOU


MariaDB
OPEN WORKS
BE UNSTOPPABLE