

# Reliability and mesh

Create resilient multi-cluster, multi-regional and multi-tenant architectures with Istio and K8s

Ameer Abbas / Google PM



#IstioCon

# Observability

**Observability**

**Security**

**Observability**

**Security**

**Traffic Mgt**

Observability

**Reliability**

Traffic Mgt

# Reliability

Architectures


# Reliability

Architectures



Multi cluster

# Reliability





# Reliability



# SLIs and SLOs

#IstioCon



## SLI

Quantitative measure of some aspect of the level of service

aka

**latency, throughput,  
availability**

## SLI

Quantitative measure of some aspect of the level of service

aka

**latency, throughput, availability**


## SLO

a target value or range of values for a service level that is measured by an SLI

aka







**99% of **Get** RPC calls will complete in less than 100 ms**

Six sided dice





Available

Outage



Available

Outage



5 / 6 availability

or

**83.33%**  
**Availability SLO**

1 / 6 outage

or

16.66%


# Aggregate SLOs

(with more dice)


#IstioCon









83.33%



83.33%




83.33%




83.33%


Outage is any (one or more) dice rolled 1




83.33%



83.33%




83.33%




83.33%

Outage is any (one or more) dice rolled 1




83.33%



83.33%



83.33%




83.33%

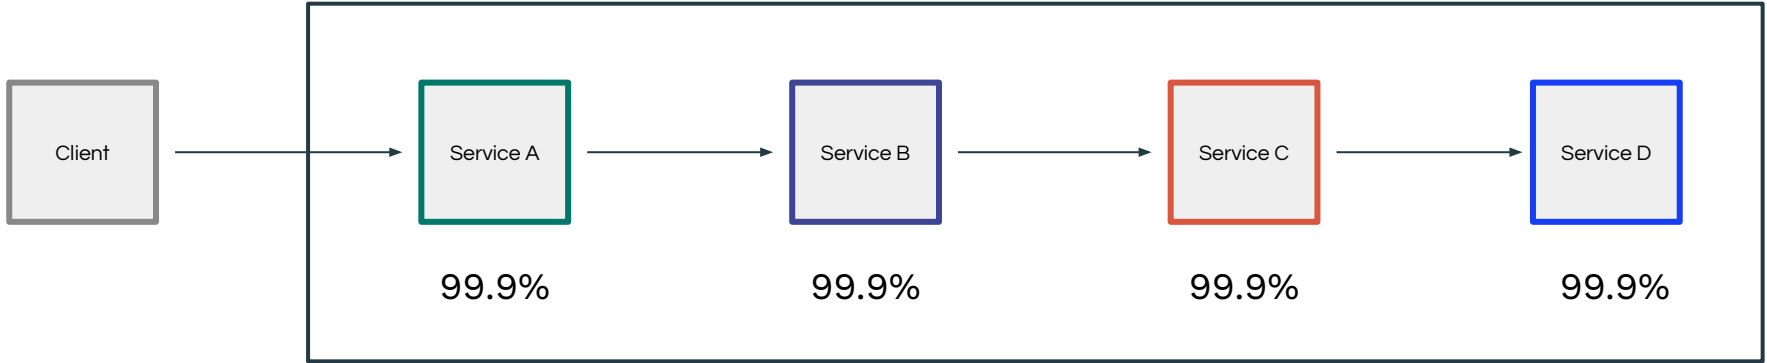
*SLO breadth*

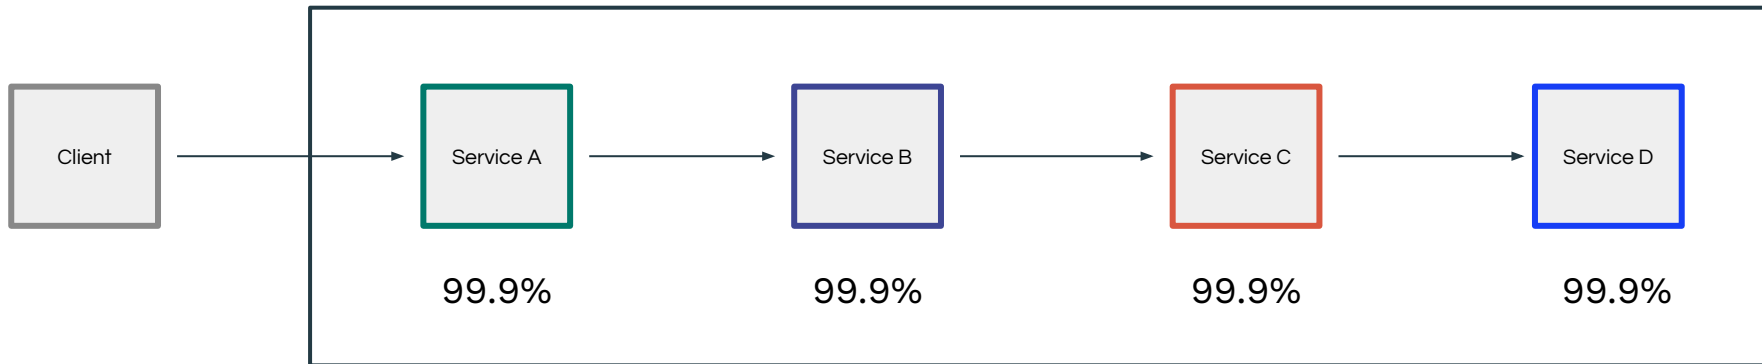
$0.8333 \times 0.8333 \times 0.8333 \times 0.8333$

**48.22% SLO**



**99.9% SLO**

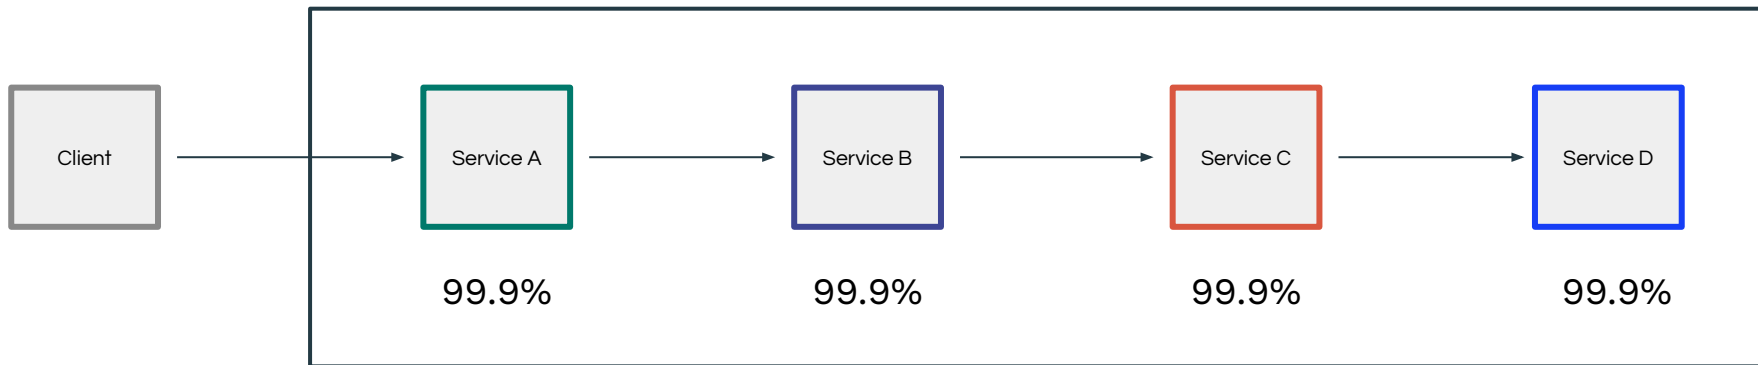




$$0.999 \times 0.999 \times 0.999 \times 0.999$$


**99.6% SLO**

## Intersection (or serial)



$$0.999 \times 0.999 \times 0.999 \times 0.999$$


**99.6% SLO**




83.33%



83.33%



83.33%




83.33%




Outage is **all dice roll 1**




83.33%



83.33%



83.33%



83.33%

Outage is **all dice roll 1**



83.33%



83.33%



83.33%

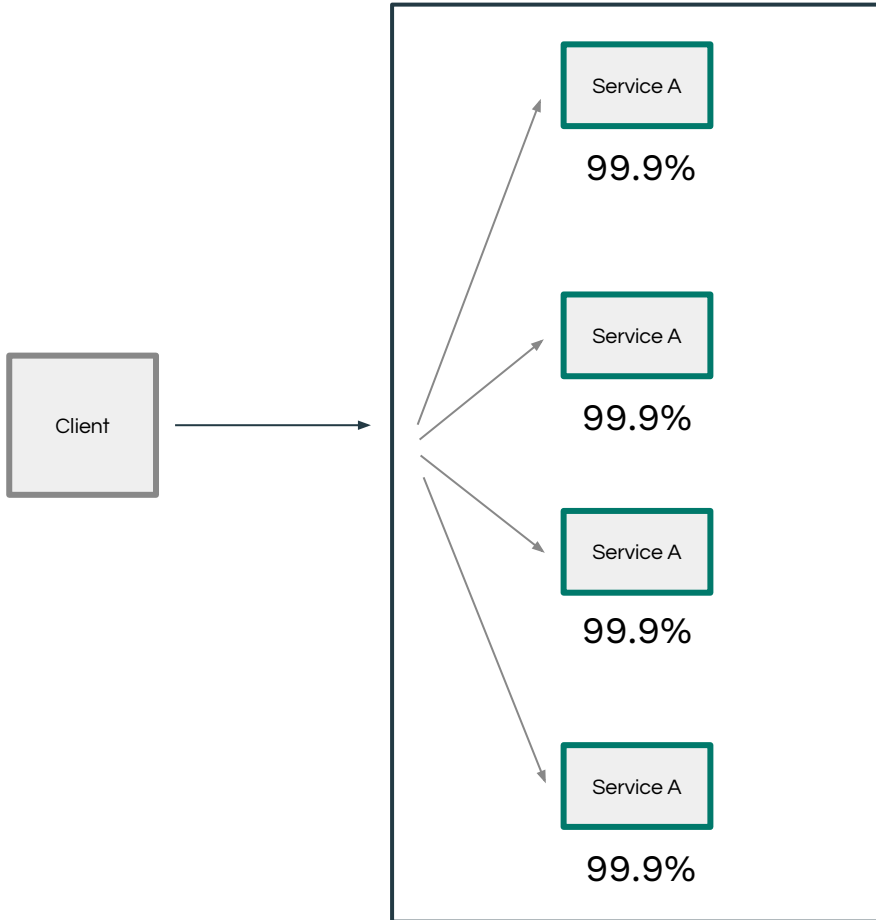


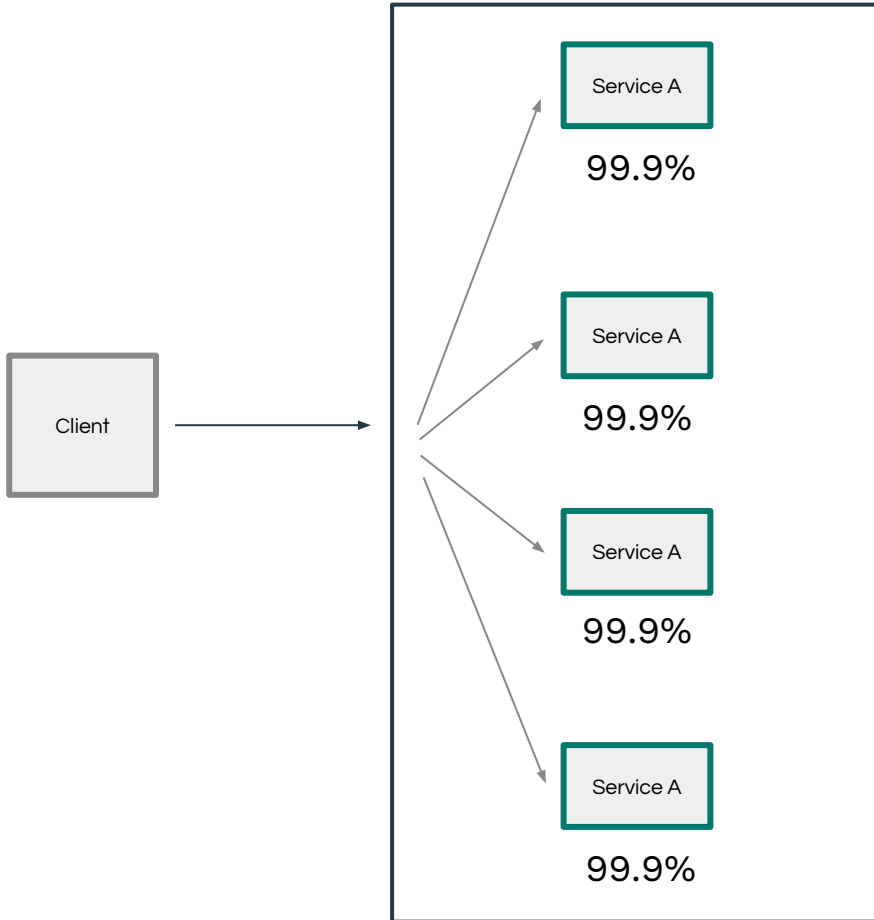
83.33%

$$1 - (\text{failure rate})^{\text{depth}}$$

$$1 - (0.1666)^4$$

**99.92% SLO**



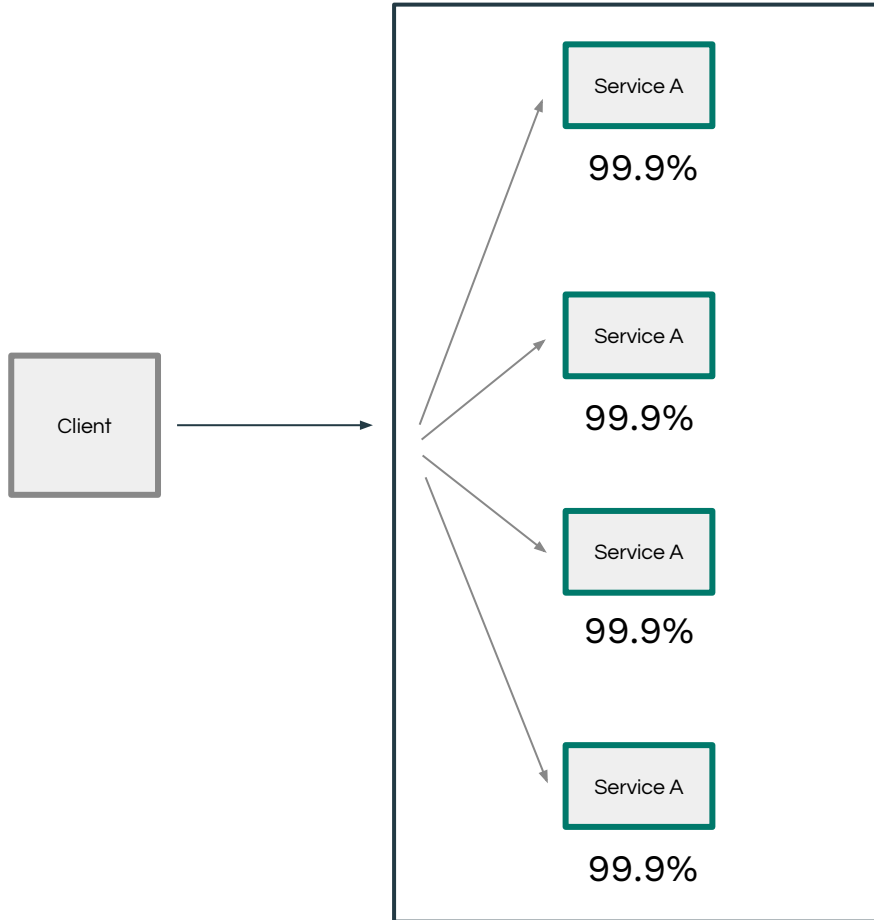


$$1 - (0.001)^4$$

**99.9999999999% SLO**

**or 11 nines**

*This is strictly mathematical and does not include any dependent variables like network, LBs, capacity planning, connectivity, and other dependent services*



**Union (aka parallel)**

$$1 - (0.001)^4$$

**99.9999999999% SLO**

**or 11 nines**

*This is strictly mathematical and does not include any dependent variables like network, LBs, capacity planning, connectivity, and other dependent services*

# Cloud Deployment Models

#IstioCon



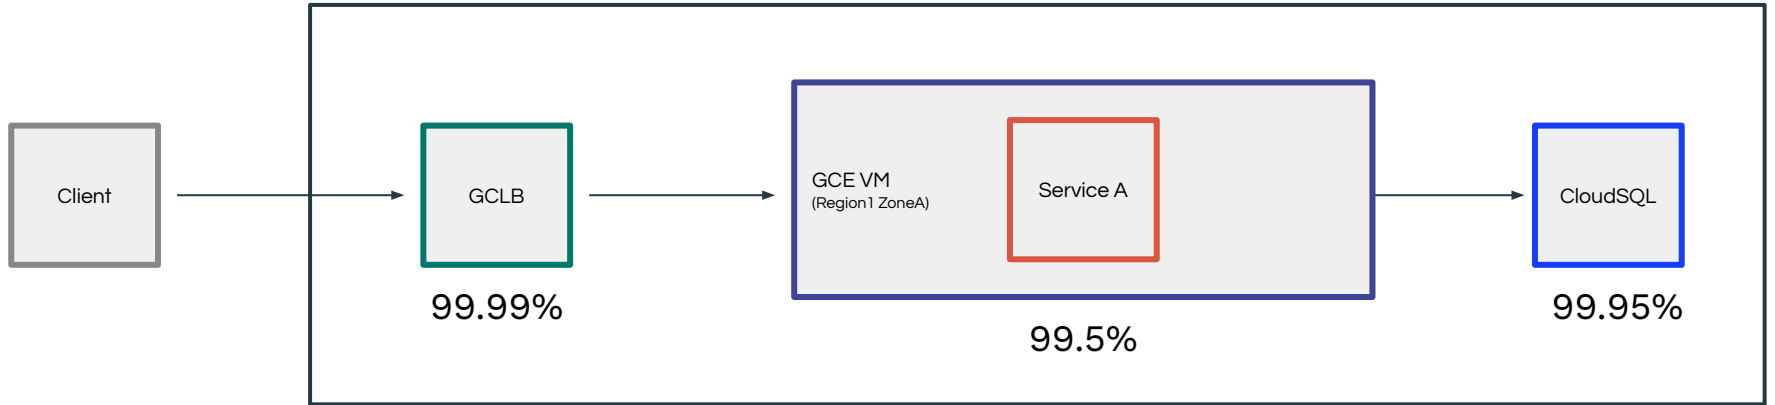
## **DISCLAIMER:**

**Following SLO calculations are mathematical and do not account for certain factors that may affect the calculation.**

**These are to be used for educational purposes only!**

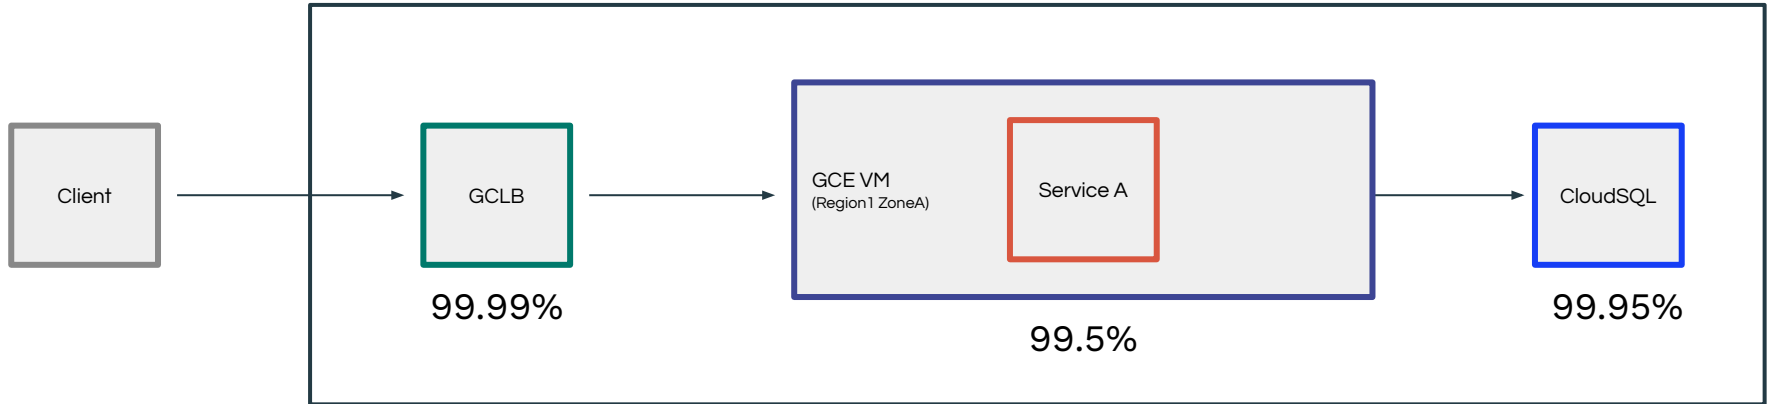


## Service in a single zonal VM





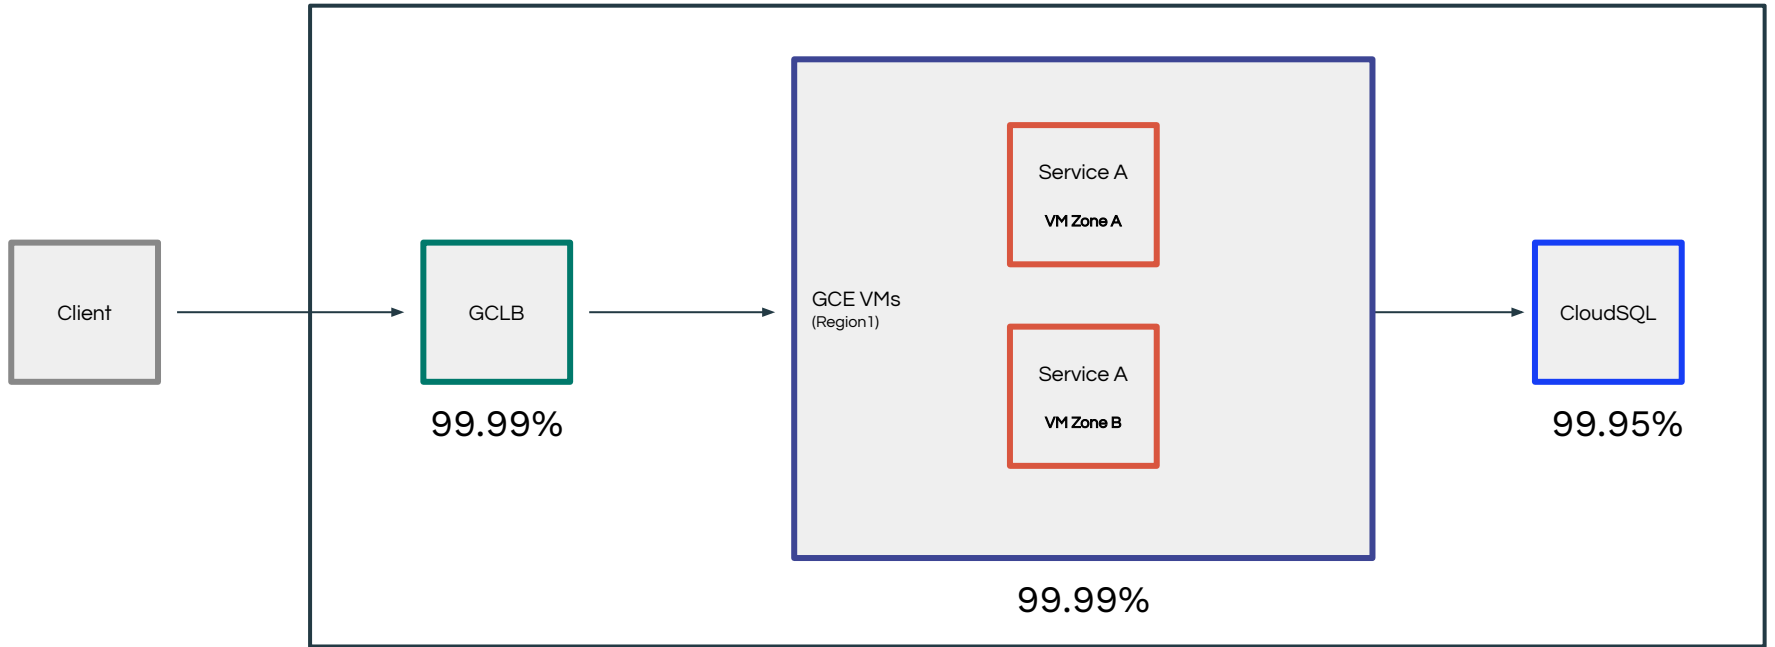
## Service in a single zonal VM



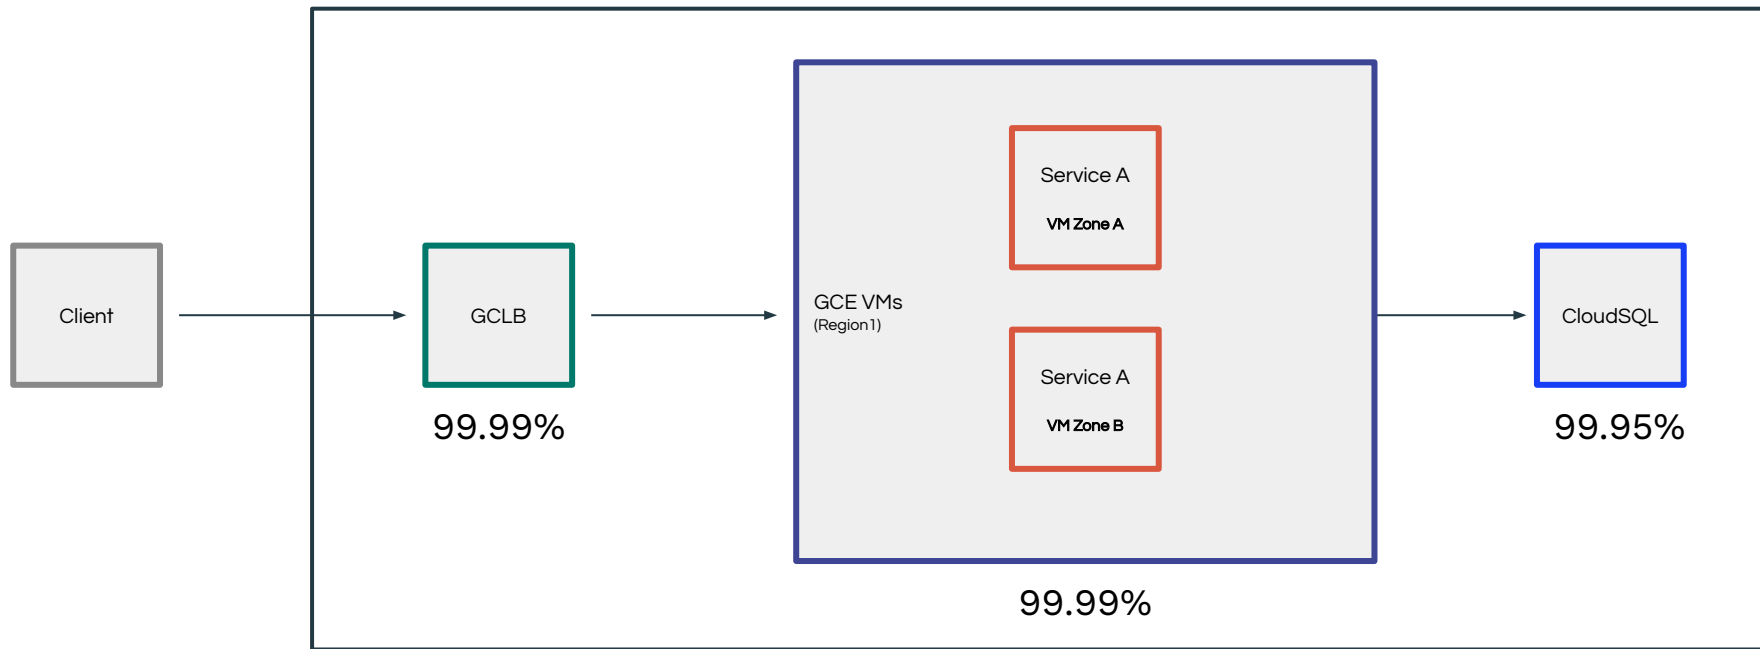
$$0.9999 \times 0.995 \times 0.9995$$

**99.44% SLO**

## Service in a multiple VMs in same region



## Service in a multiple VMs in same region



$$0.9999 \times 0.9999 \times 0.9995$$

**99.93% SLO**

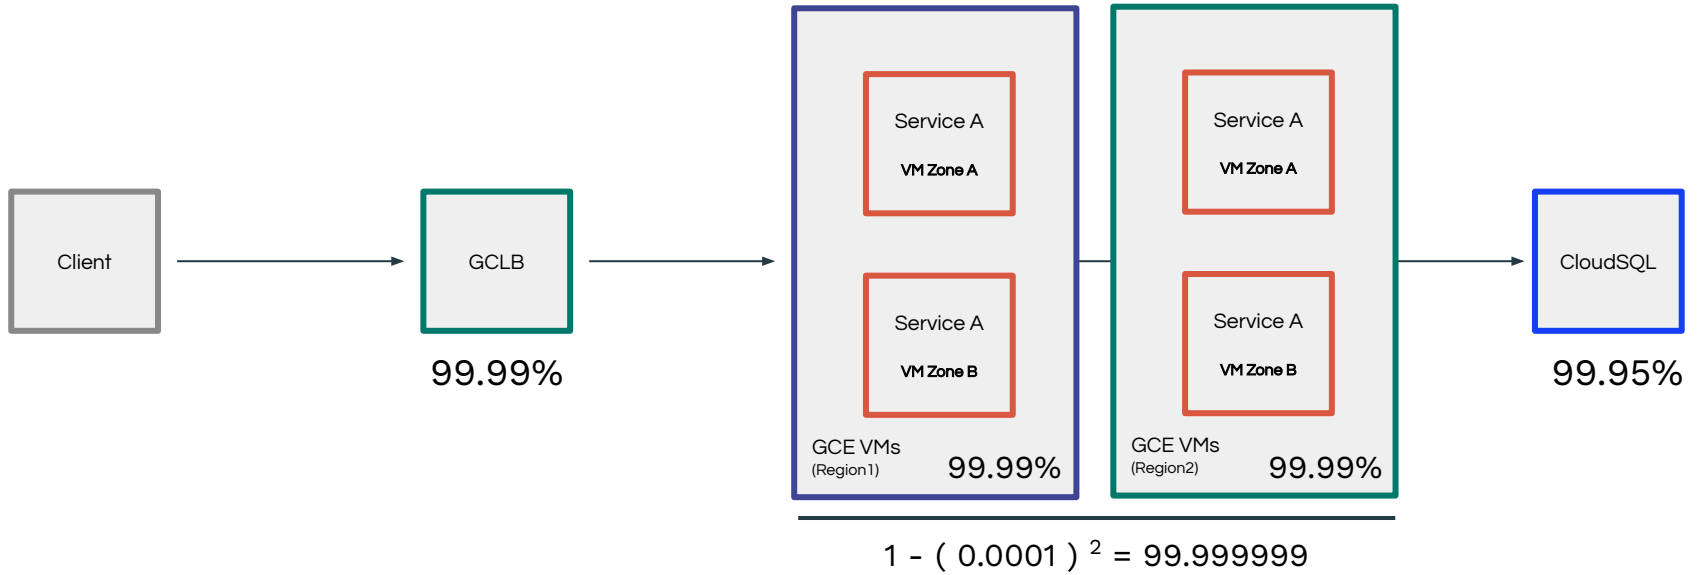
Multiple zonal VMs in 1 Region	<b>99.93%</b>
Single zonal VM	99.44%

## Service in a multiple VMs in multiple regions



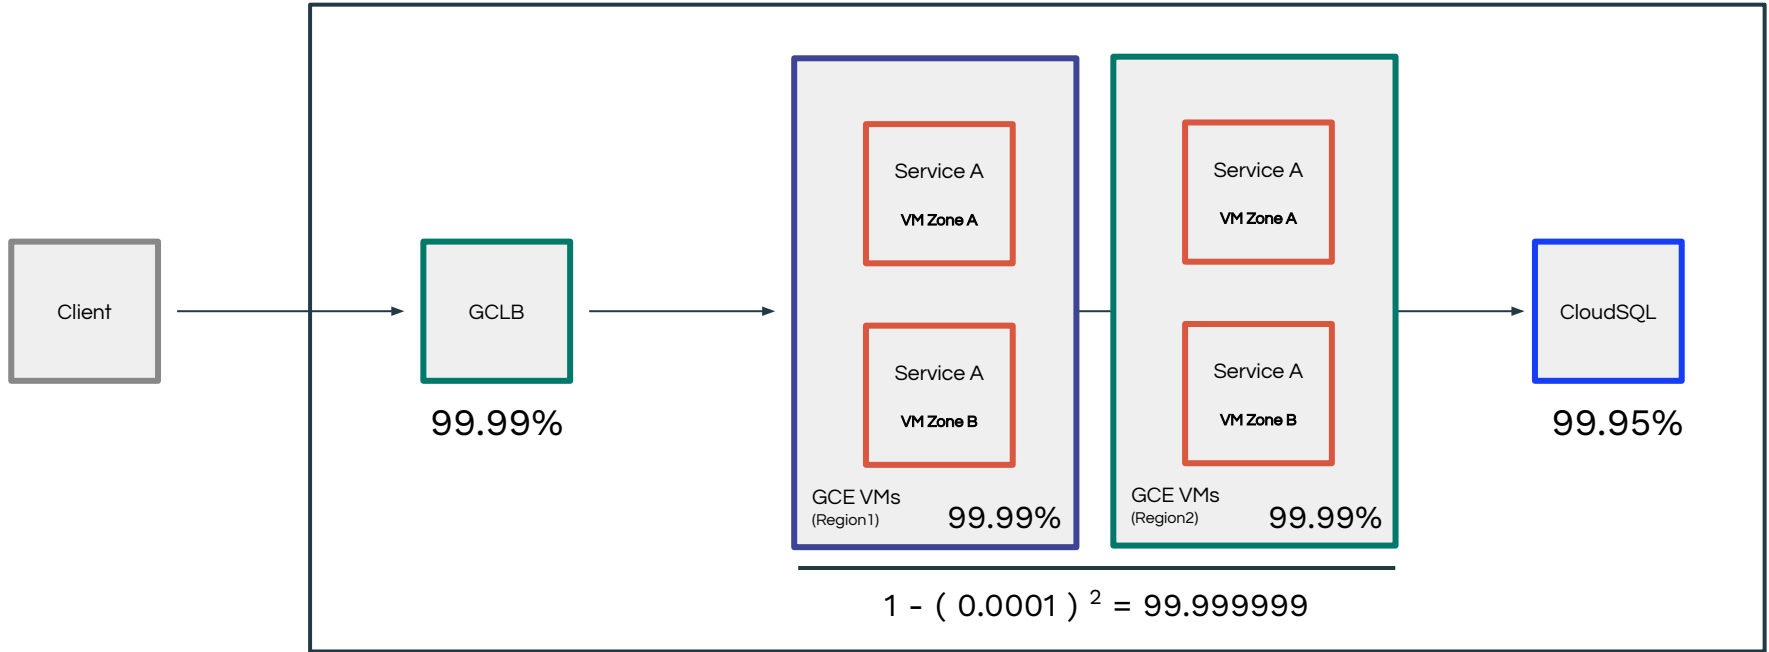
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in a multiple VMs in multiple regions



Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in a multiple VMs in multiple regions



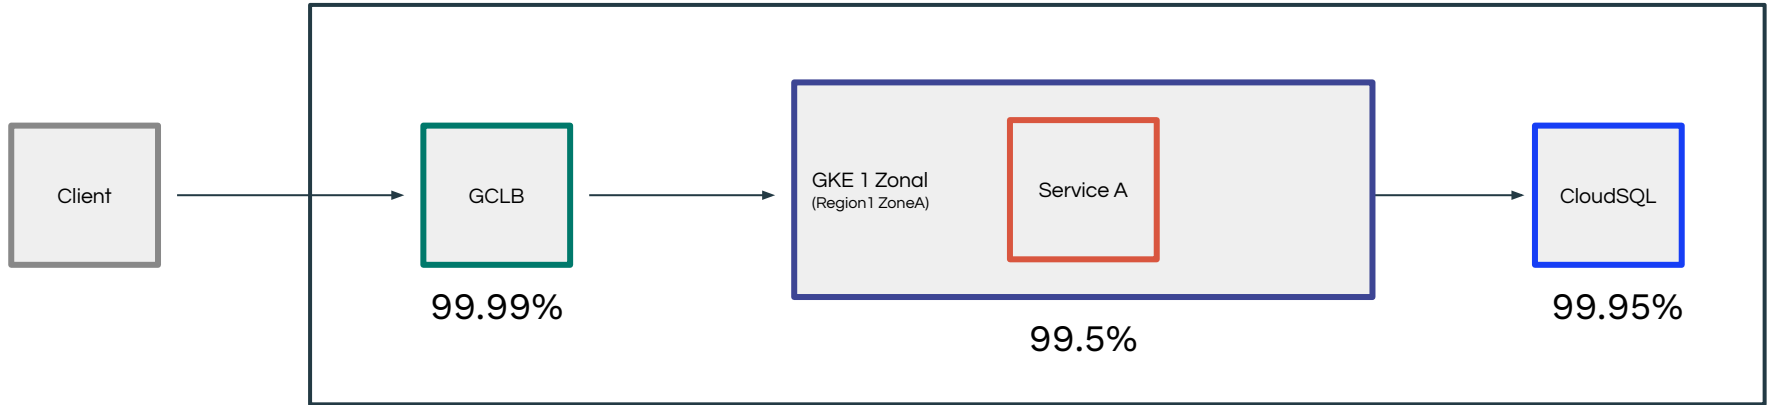
$$0.9999 \times 0.999999999 \times 0.9995$$

**99.94% SLO**

Multi regional multi zonal VMs	<b>99.94%</b>
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

**enough about VMs**

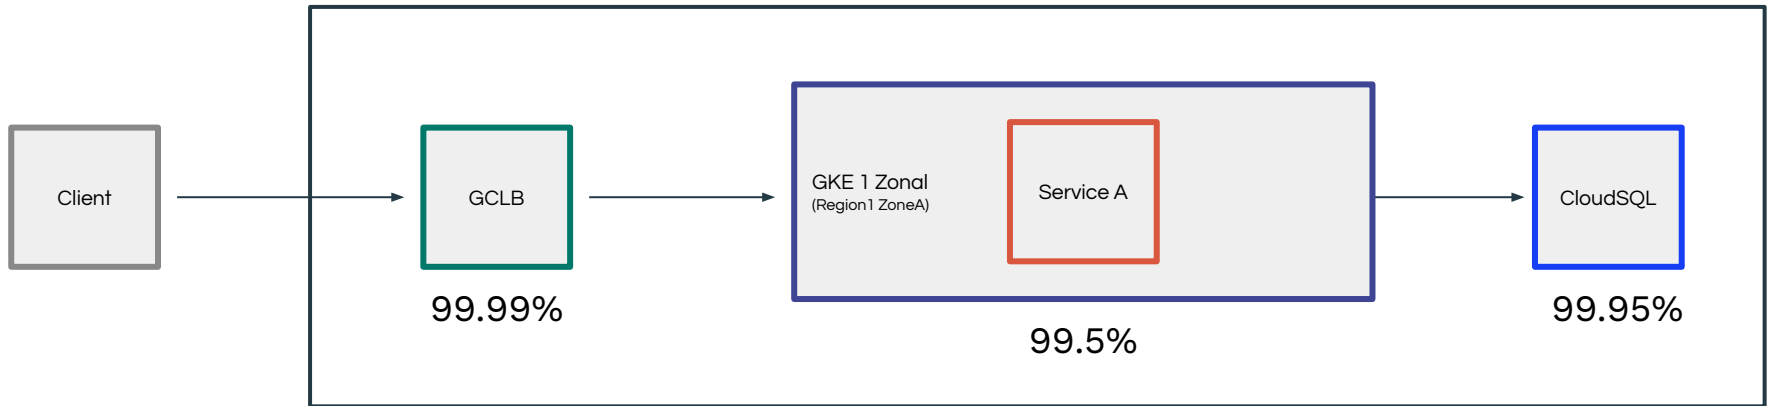
## Service in a single zonal GKE cluster



Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%



## Service in a single zonal GKE cluster




$$0.9999 \times 0.995 \times 0.9995$$

**99.44% SLO**

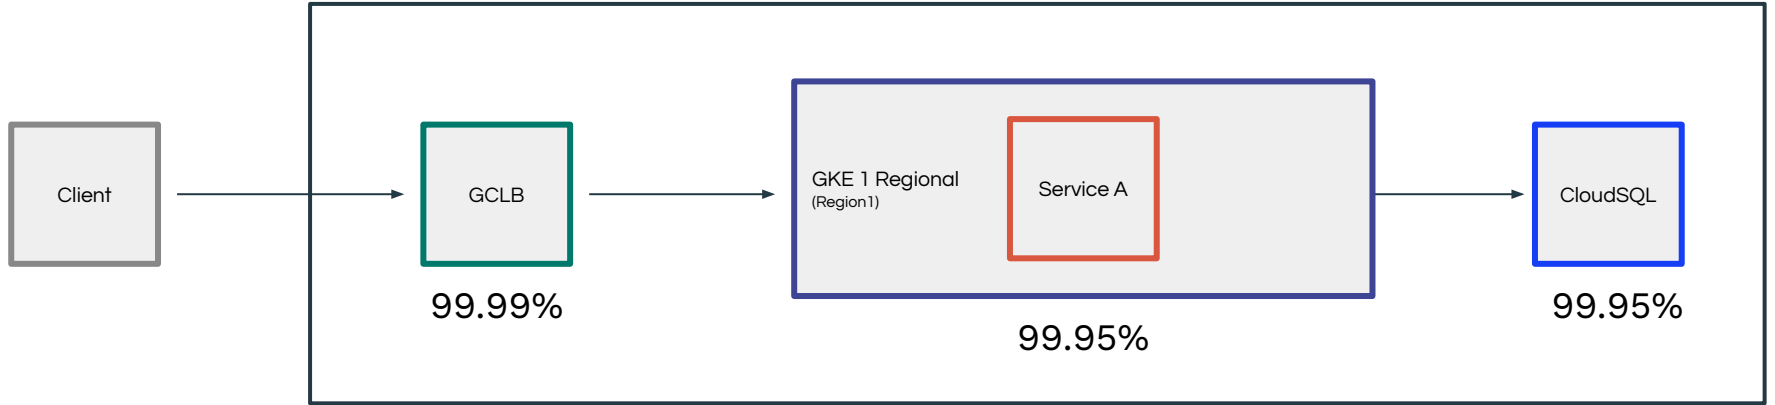
Single zonal K8s cluster	<b>99.44%</b>
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in a single regional GKE cluster



Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in a single regional GKE cluster

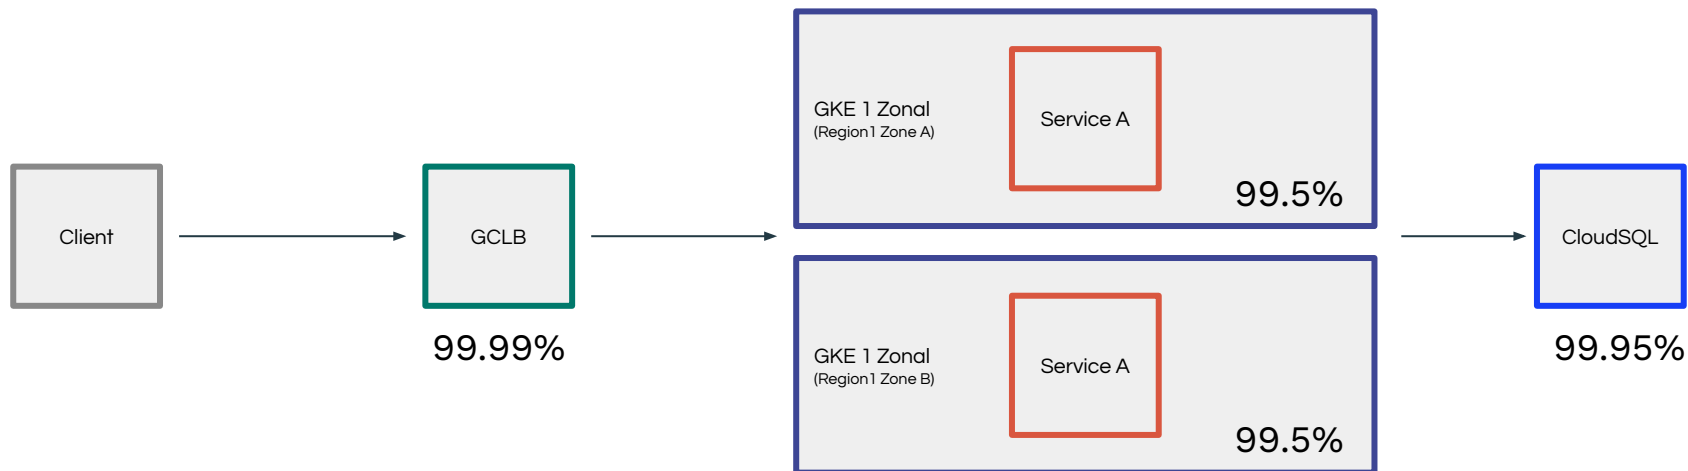


$$0.9999 \times 0.9995 \times 0.9995$$

**99.89% SLO**

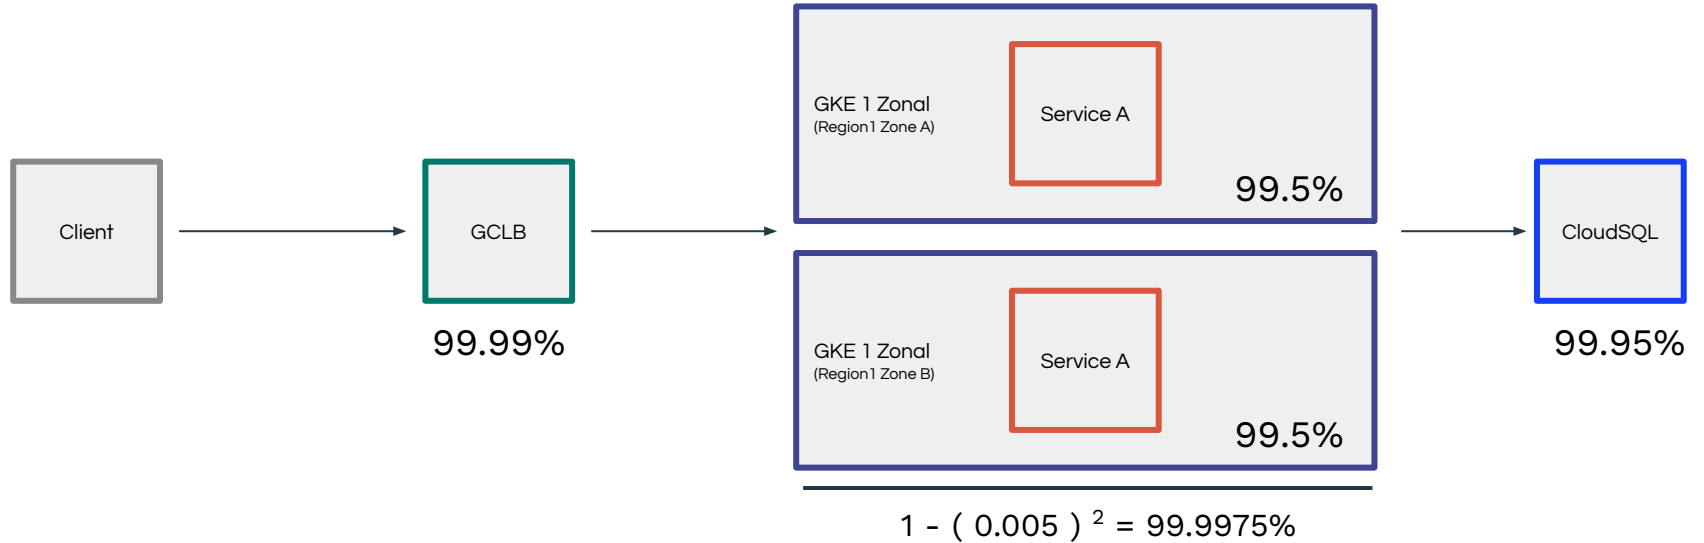
Single Regional K8s cluster	<b>99.89%</b>
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in multiple zonal GKE clusters in one region



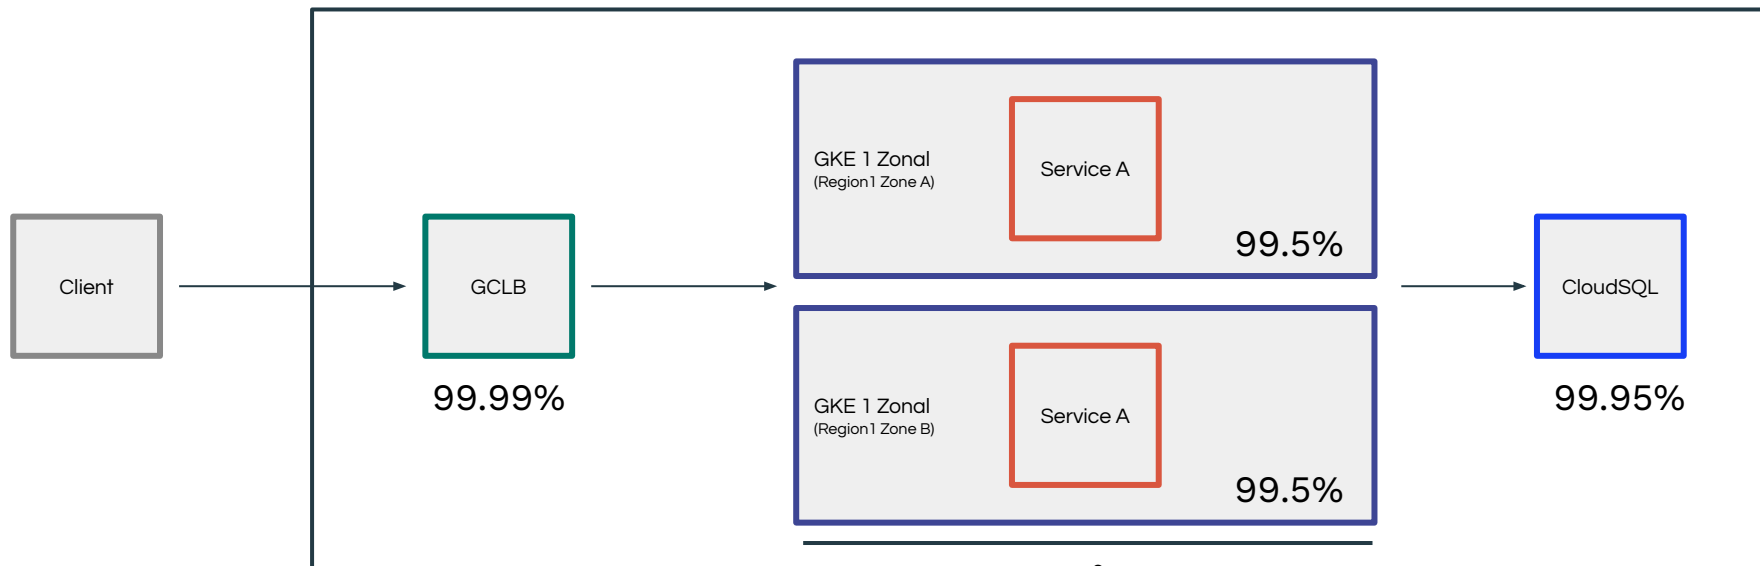
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in multiple zonal GKE clusters in one region



Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in multiple zonal GKE clusters in one region



Multiple zonal K8s cluster 1 region	<b>99.93%</b>
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

$$1 - (0.005)^2 = 99.9975\%$$

$$0.9999 \times 0.999975 \times 0.9995$$

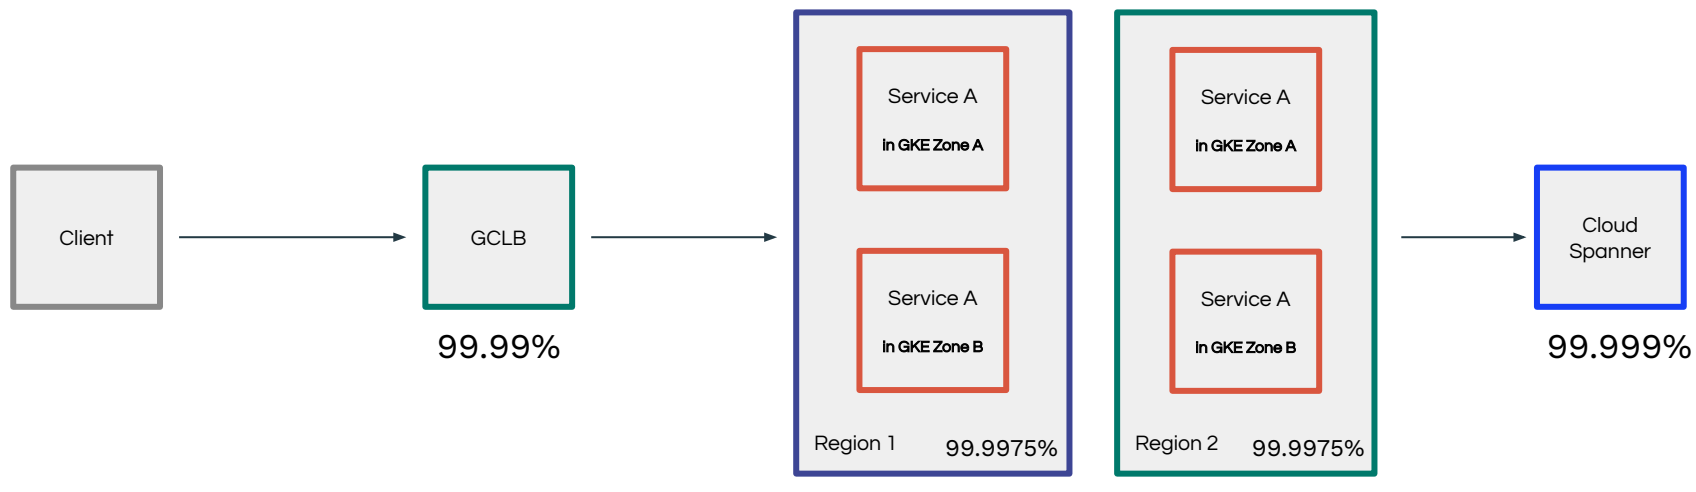
**99.93% SLO**

## Service in multiple zonal GKE clusters in multiple regions



Multiple zonal K8s cluster 1 region	99.93%
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## Service in multiple zonal GKE clusters in multiple regions

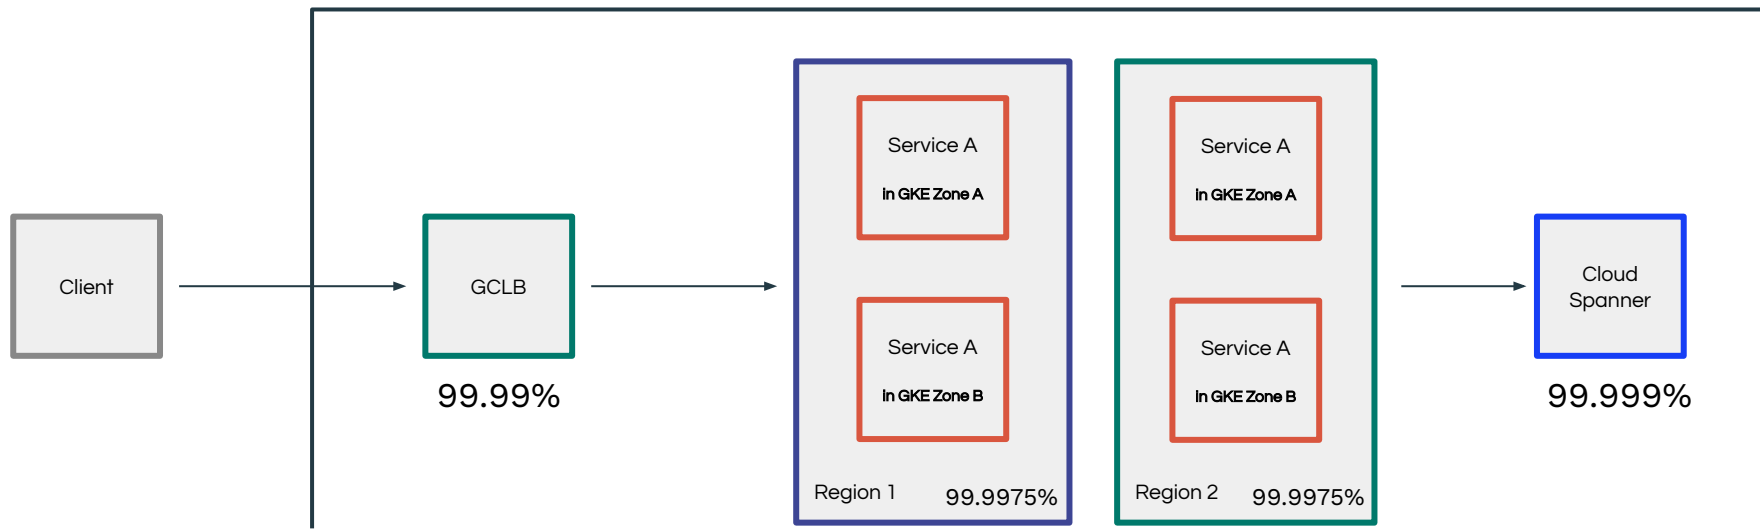


$$1 - (0.000025)^2 = 99.9999999\%$$

Multiple zonal K8s cluster 1 region	99.93%
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%



## Service in multiple zonal GKE clusters in multiple regions



$$1 - (0.000025)^2 = 99.99999999\%$$

$$0.9999 \times 0.9999999999 \times 0.99999$$

**99.989% SLO**

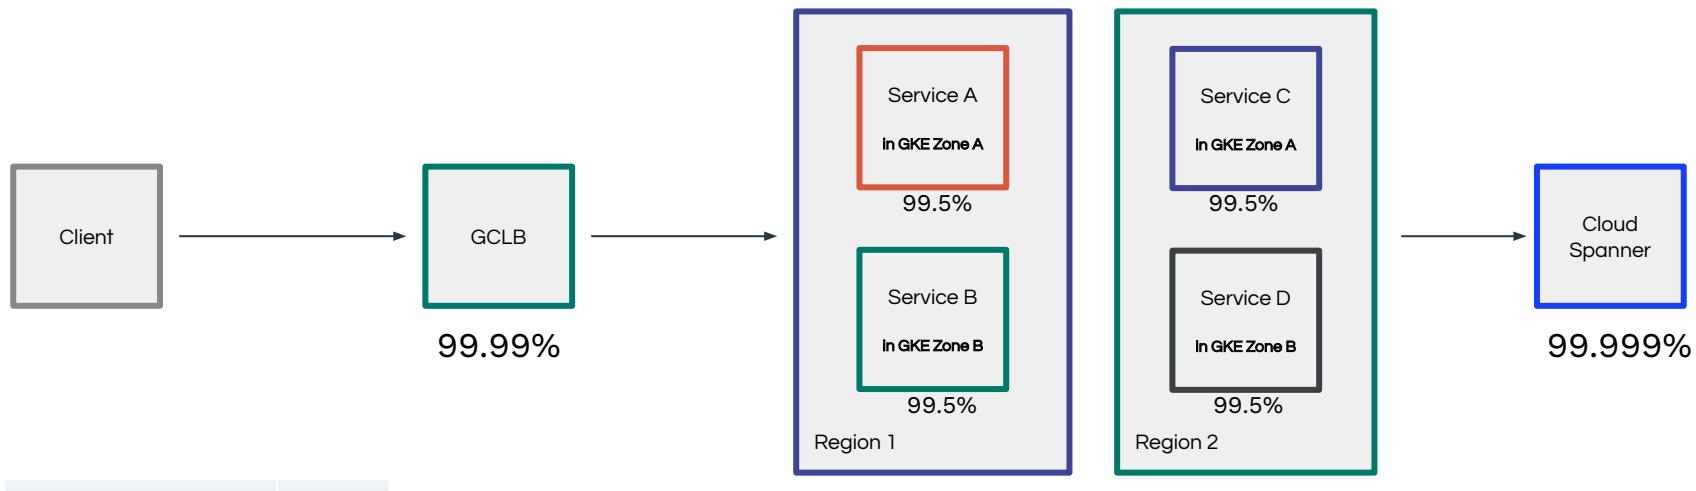
Multiple zonal and regional K8s	<b>99.99%</b>
Multiple zonal K8s cluster 1 region	99.93%
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

## One tenant (service) per GKE cluster



Multiple zonal and regional K8s	99.99%
Multiple zonal K8s cluster 1 region	99.93%
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

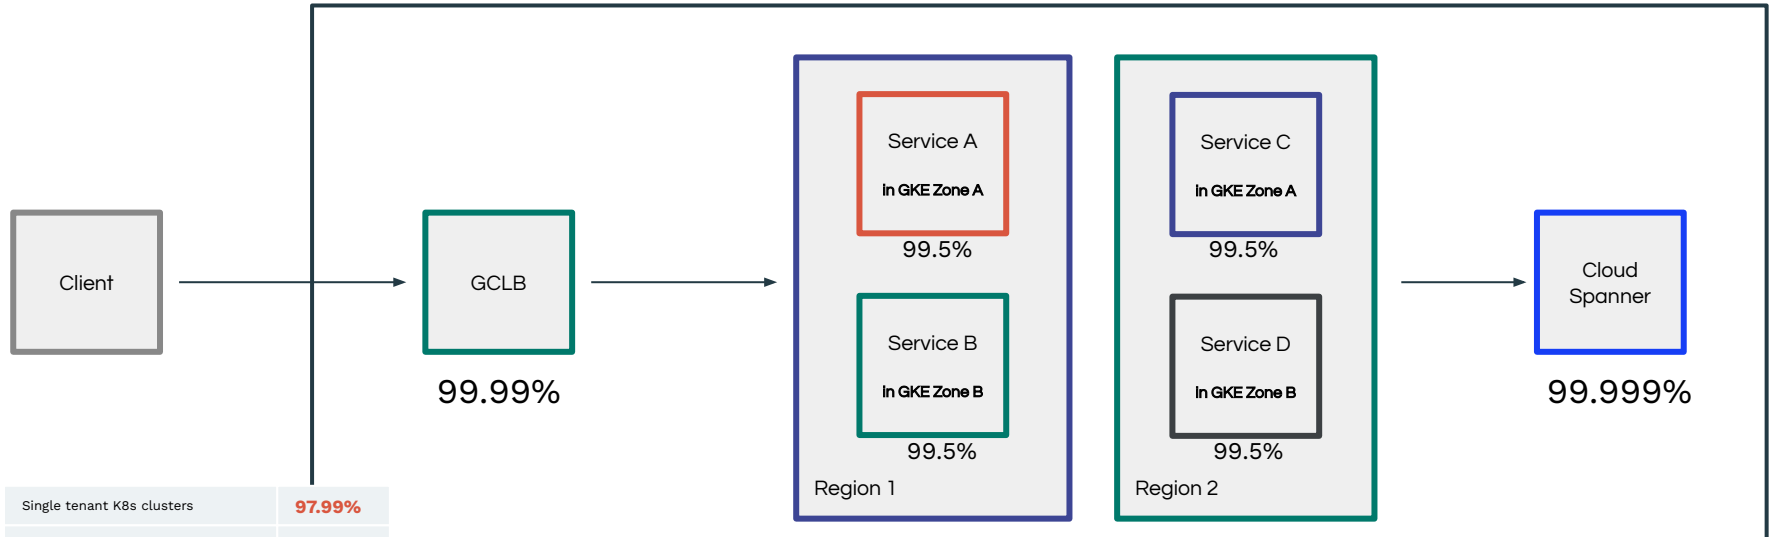
## One tenant (service) per GKE cluster



$$0.995 \times 0.995 \times 0.995 \times 0.995 = 98.01\%$$

Multiple zonal and regional K8s	99.99%
Multiple zonal K8s cluster 1 region	99.93%
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%

# One tenant (service) per GKE cluster

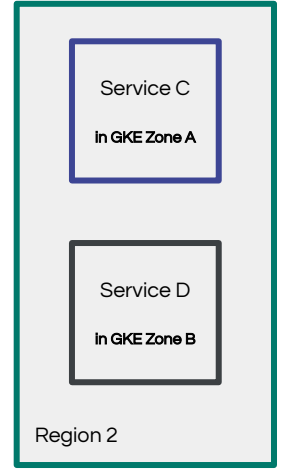

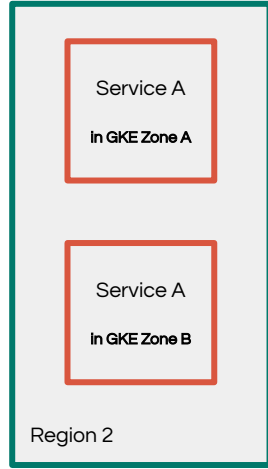



$$0.995 \times 0.995 \times 0.995 \times 0.995 = 98.01\%$$

$$0.9999 \times 0.9801 \times 0.99999$$

**97.99% SLO**

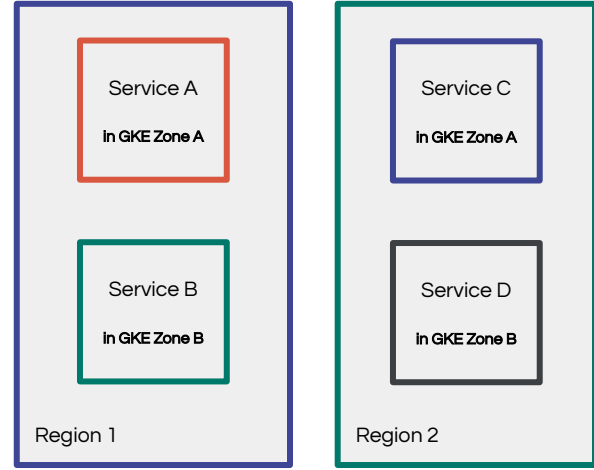
Single tenant K8s clusters	<b>97.99%</b>
Multiple zonal and regional K8s	<b>99.99%</b>
Multiple zonal K8s cluster 1 region	99.93%
Single Regional K8s cluster	99.89%
Single zonal K8s cluster	99.44%
Multi regional multi zonal VMs	99.94%
Multiple zonal VMs in 1 Region	99.93%
Single zonal VM	99.44%





---

$$1 - (0.000025)^2 = \mathbf{99.999999\%}$$



---


$$0.995 \times 0.995 \times 0.995 \times 0.995 = \mathbf{98.01\%}$$



---


$$1 - ( 0.000025 ) ^ 2 = \mathbf{99.999999\%}$$

**31.56 milliseconds of yearly downtime**



---


$$0.995 \times 0.995 \times 0.995 \times 0.995 = \mathbf{98.01\%}$$

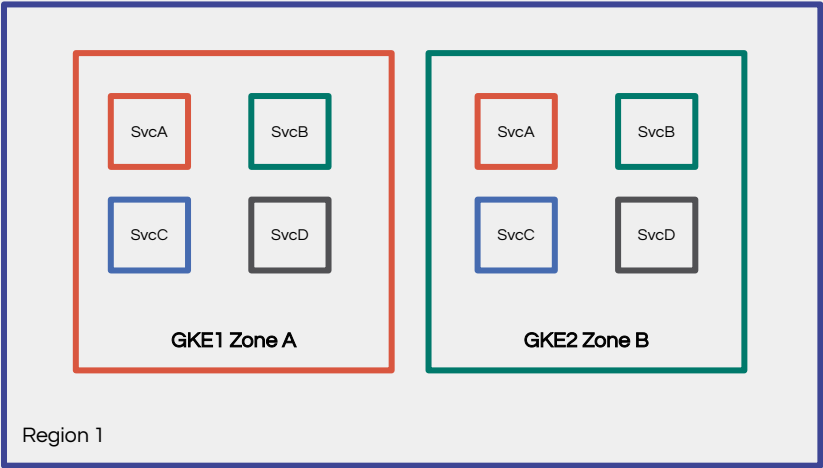
**14 hours 32 minutes of monthly downtime**

# Multi regional Distributed Services





# Mesh and Distributed Services



Mesh

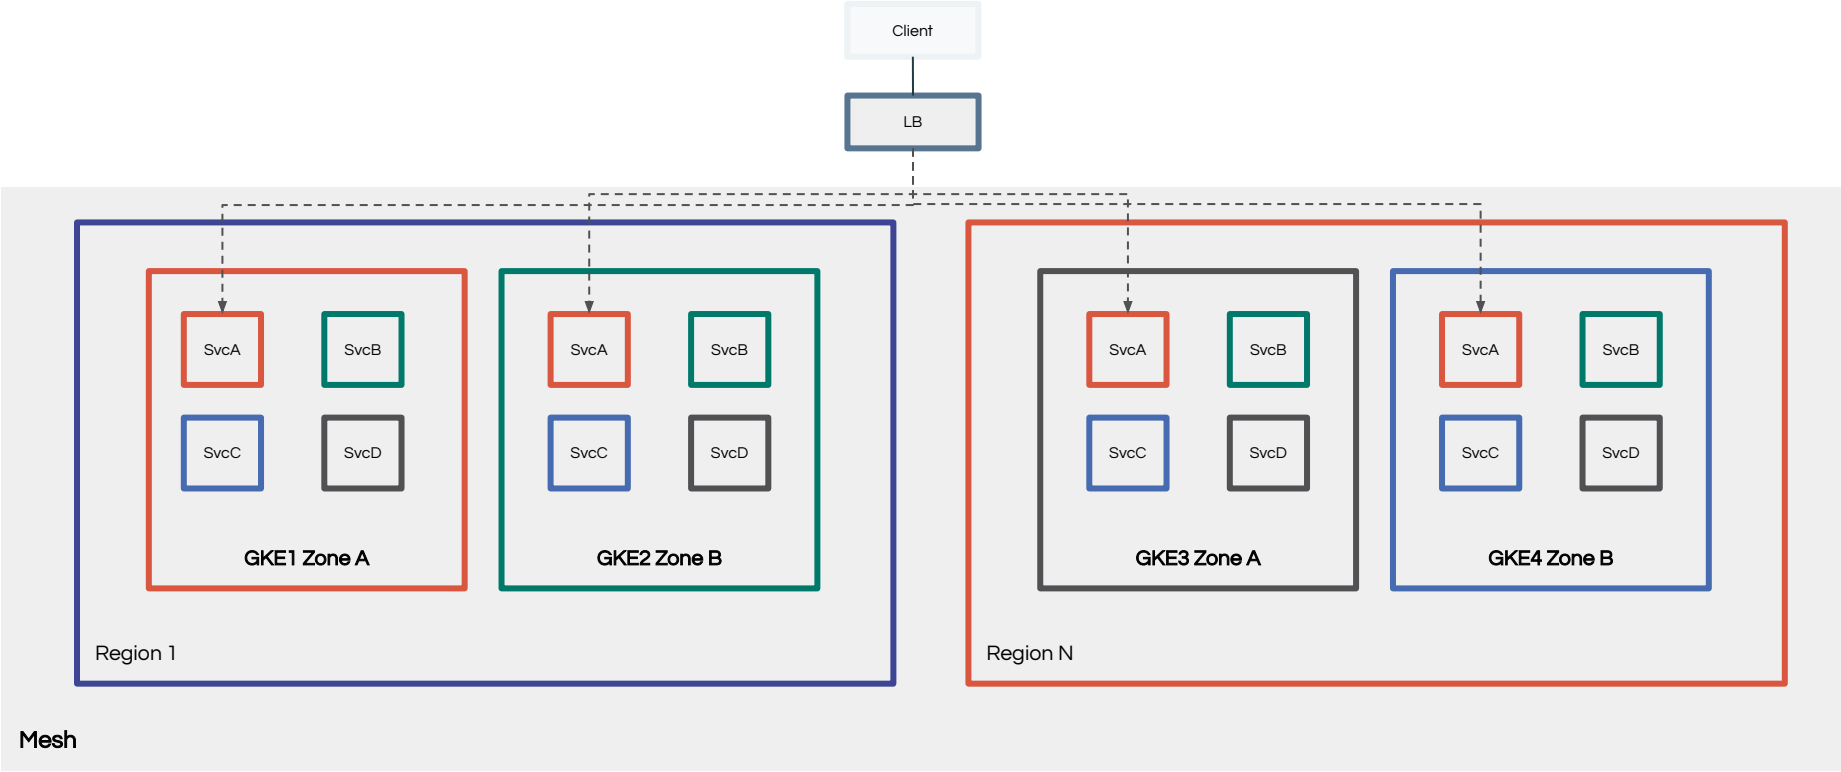
# Mesh and Multiregional Distributed Services

#IstioCon




# Locality considerations

# Locality considerations





# Locality considerations

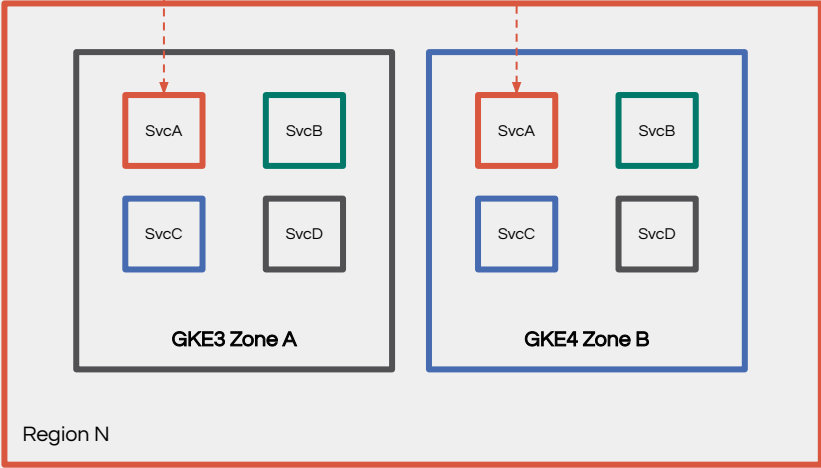
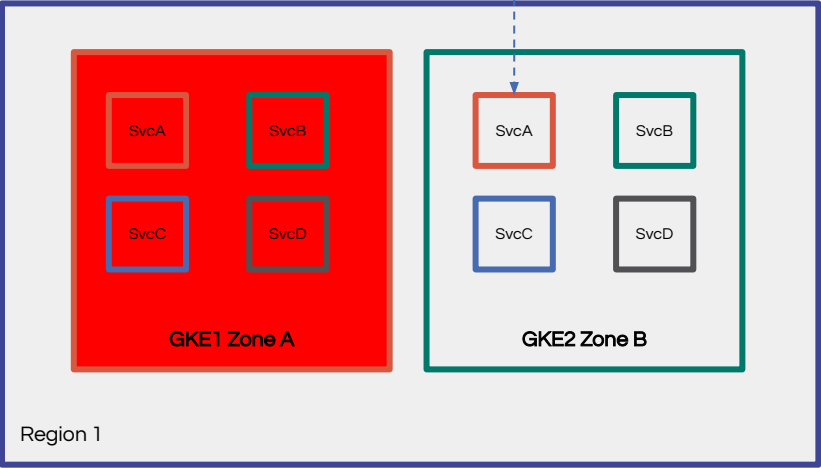



# **Business continuity**

# Business continuity

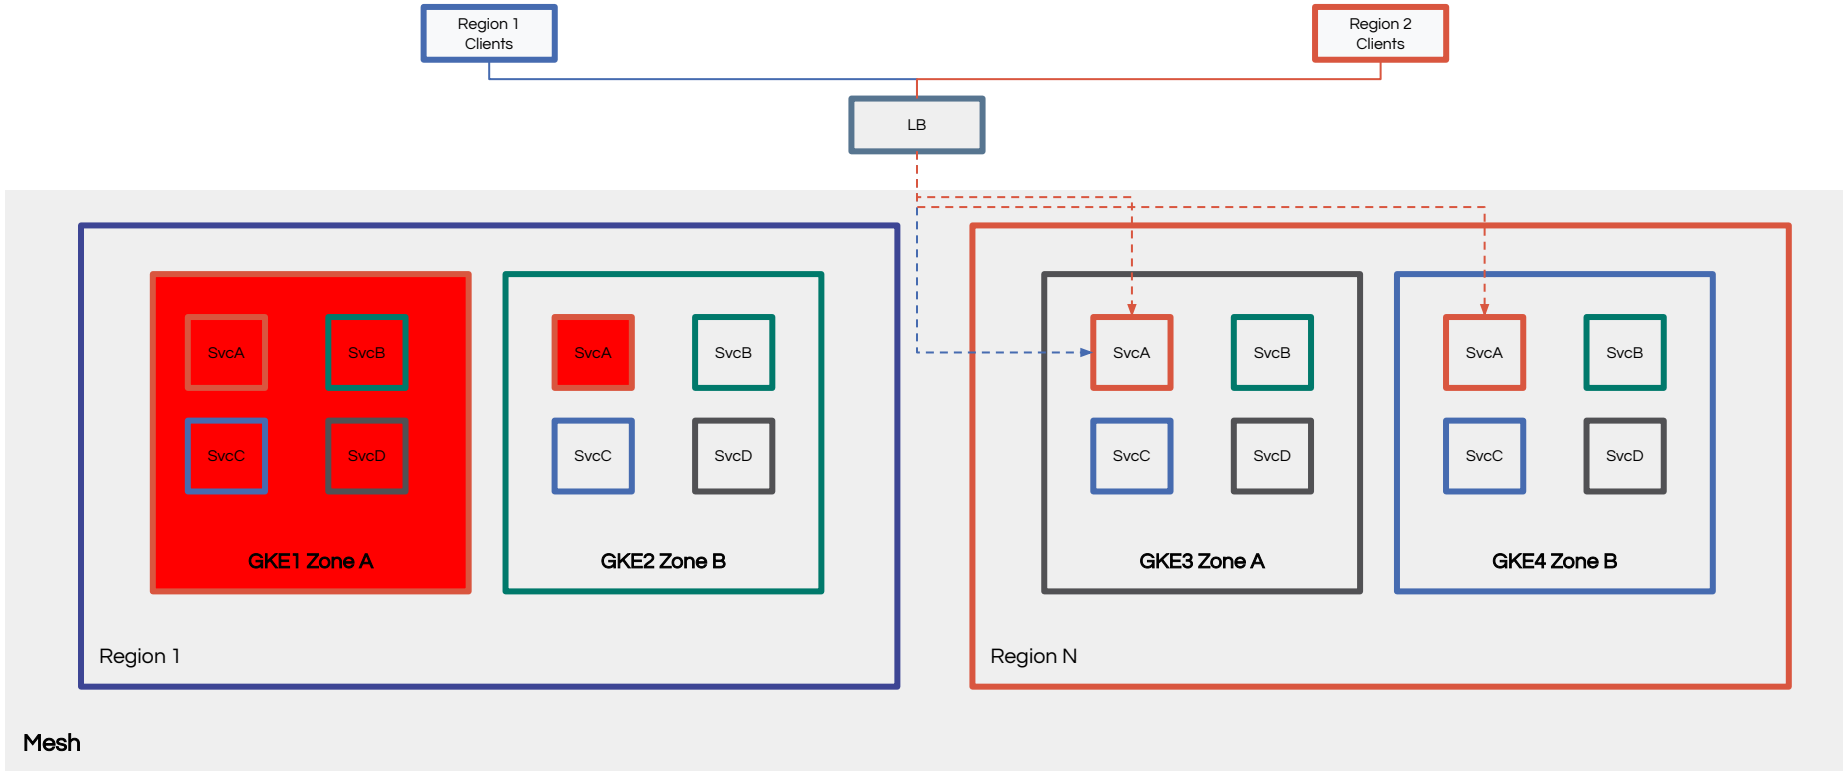


# Business continuity

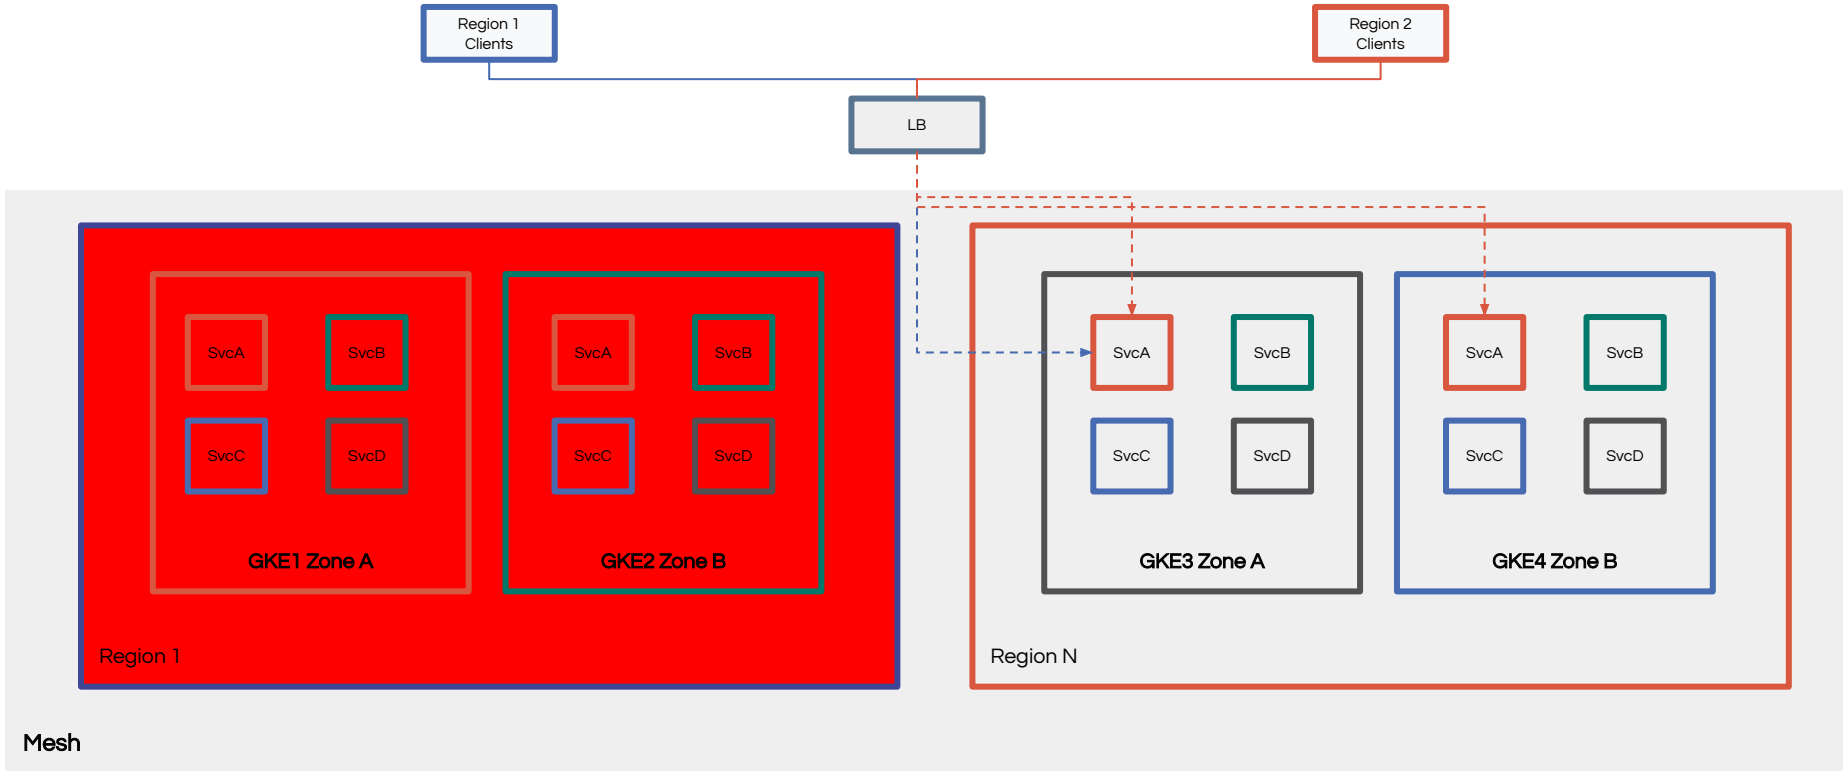




# Business continuity

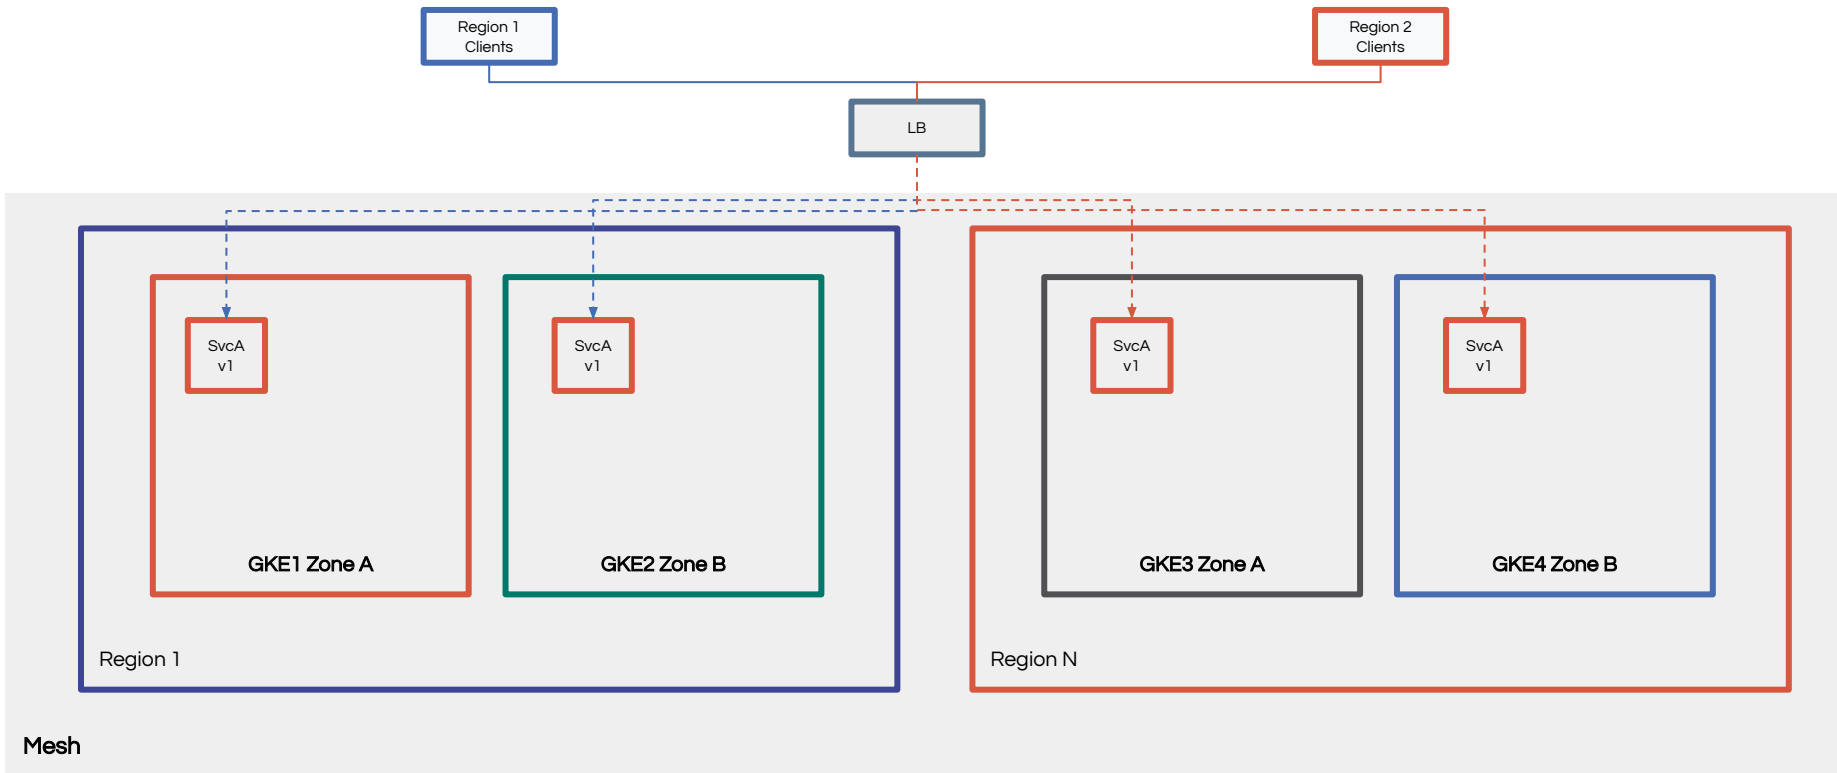


# Business continuity

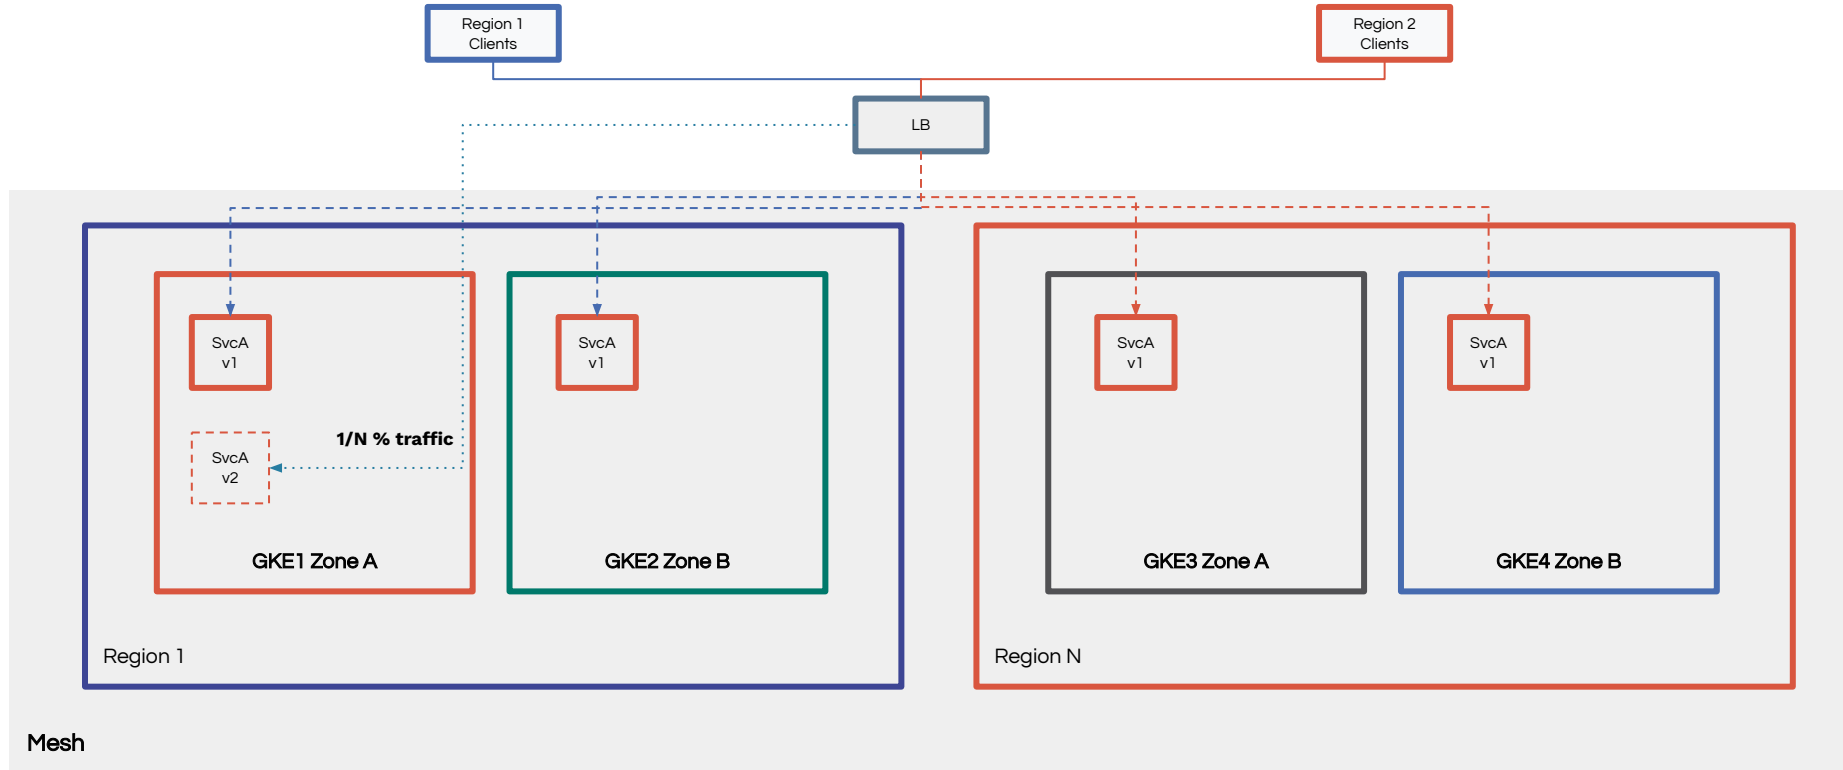


# **Application lifecycle management**

# Application lifecycle


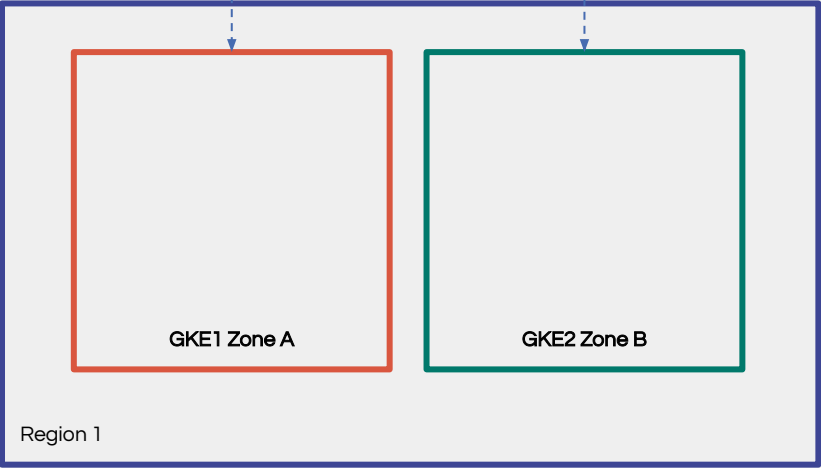



# Application lifecycle



# **Infrastructure lifecycle management**

# Infrastructure lifecycle




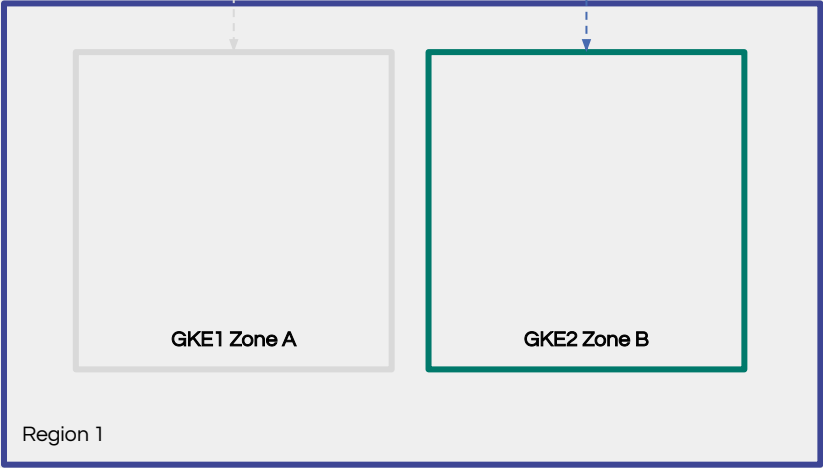
# Infrastructure lifecycle

Region 1  
Clients

Region 2  
Clients


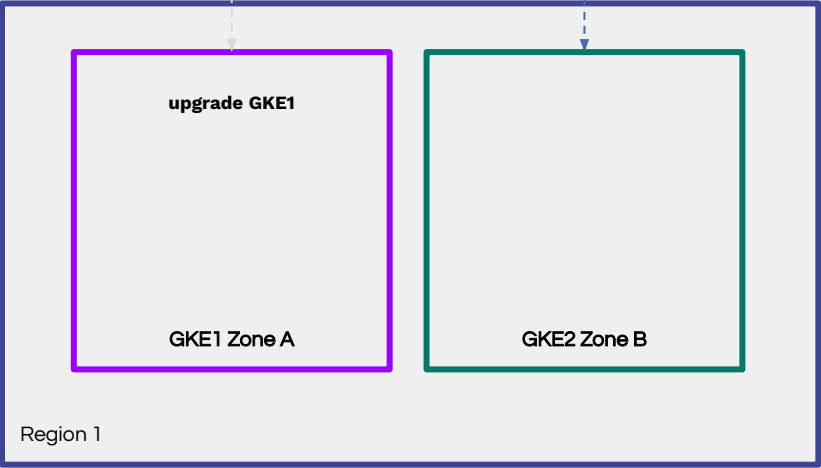

LB

drain GKE1


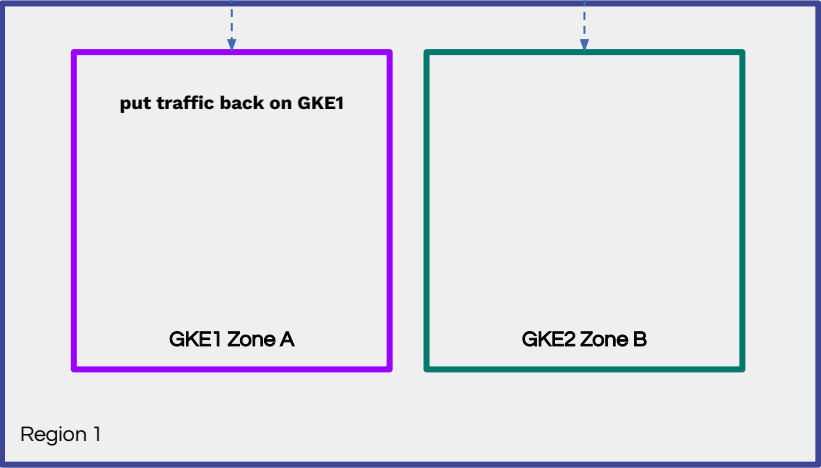






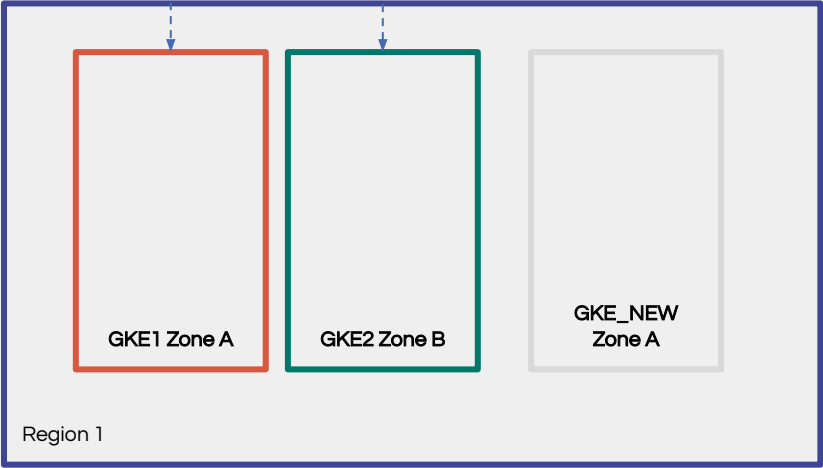

# Infrastructure lifecycle




# Infrastructure lifecycle




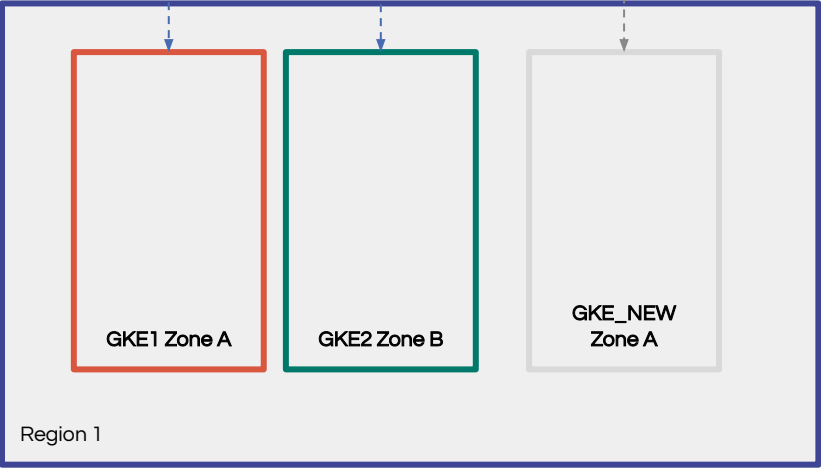
# Infrastructure lifecycle



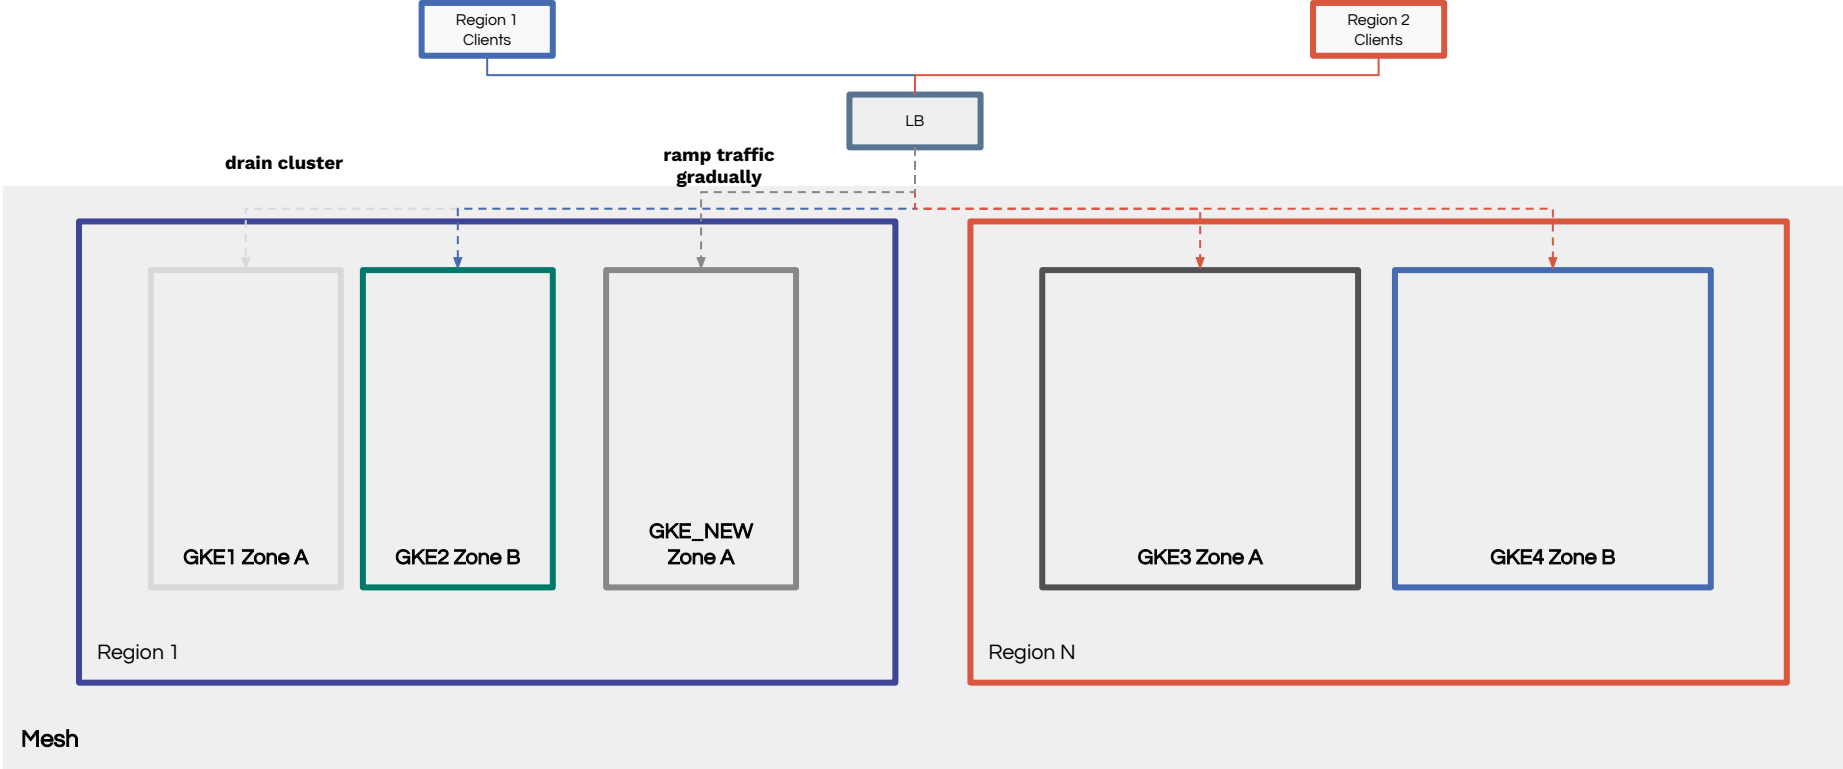
# Infrastructure lifecycle



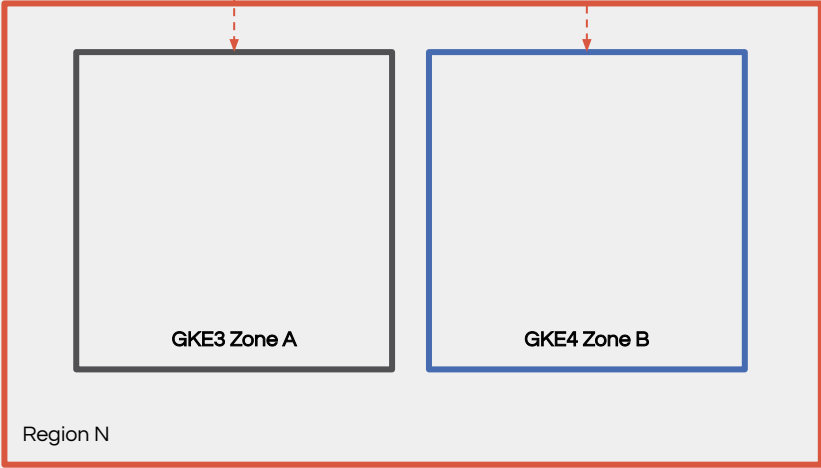
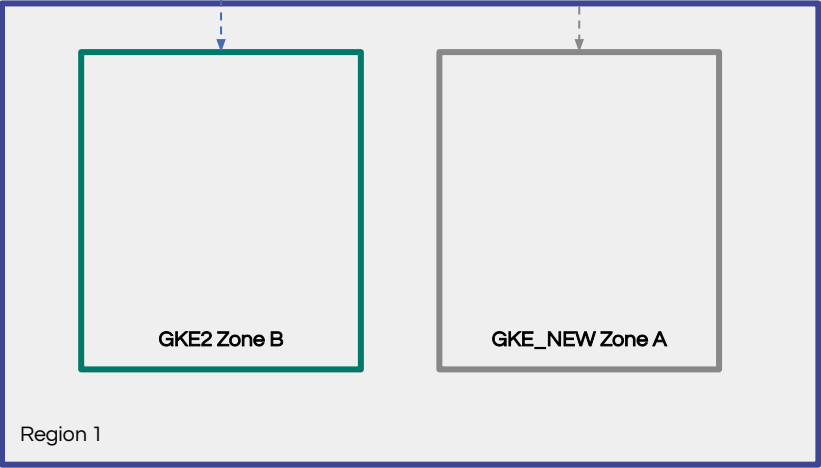

ramp traffic gradually



# Infrastructure lifecycle

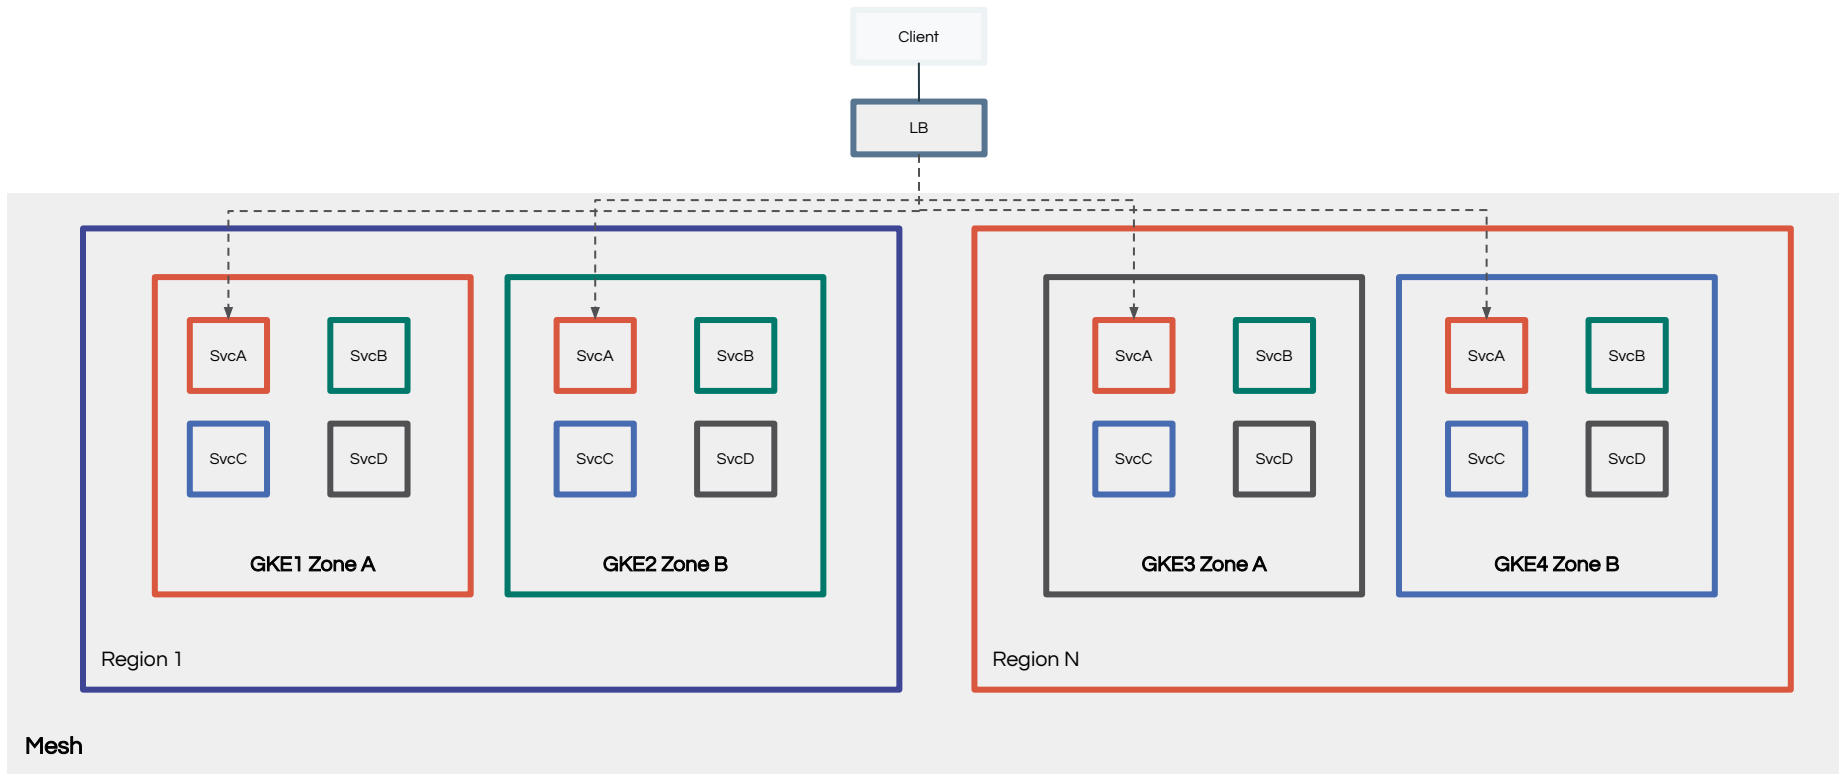


# Infrastructure lifecycle



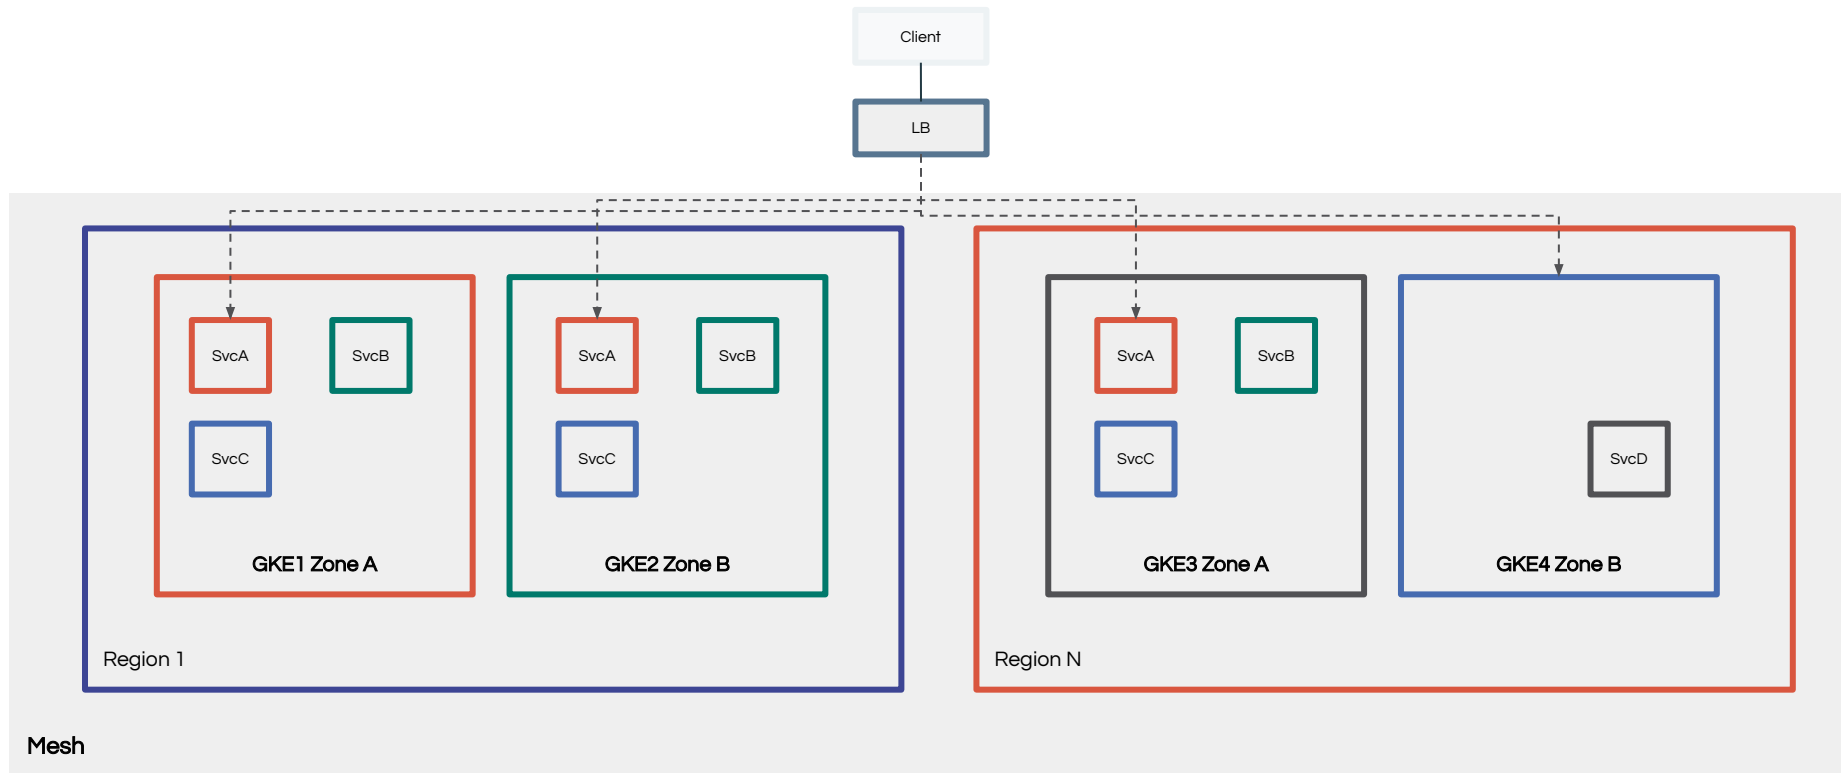
**Single tenant clusters?**

# Tenancy

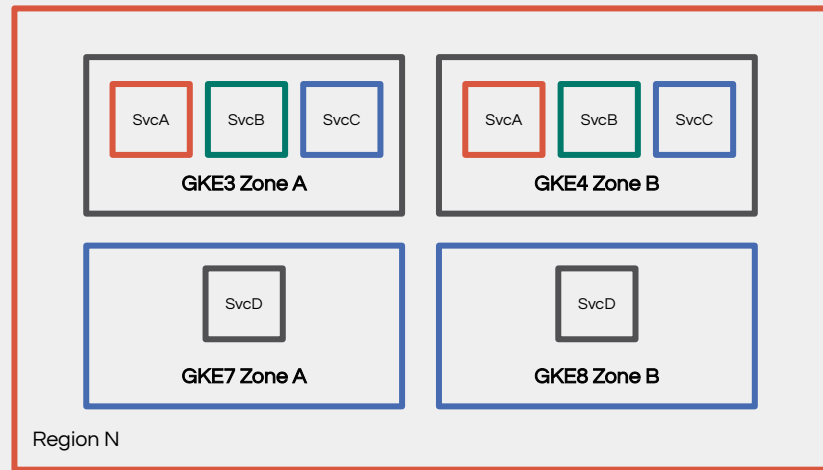
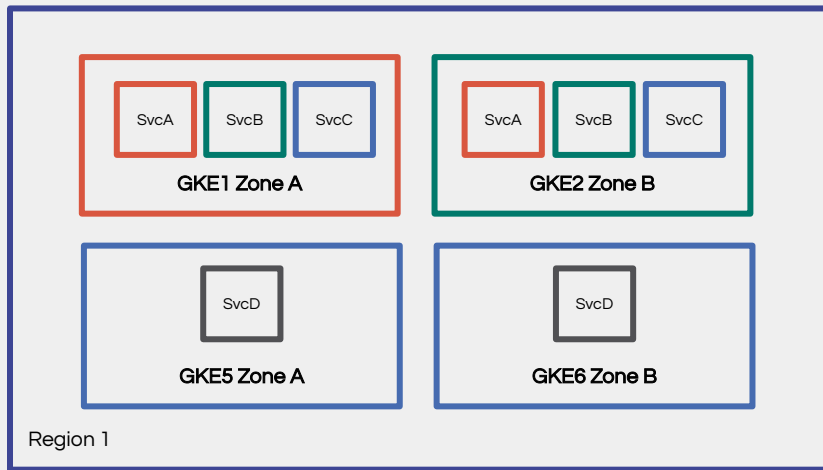





# Tenancy



# Tenancy



# Final thoughts

On Reliability

On Application management

On Infrastructure management

On Tenancy

#IstioCon



# Thank you!

#IstioCon

