

Kubernetes on EDGE DAY

NORTH AMERICA

Kubernetes on EDGE DAY
NORTH AMERICA

When the Region Goes Dark

Kubernetes-Native Edge Failover with k8gb



Fabrizio Sgura
Chief Engineer, Veritas Automata
CNCF Ambassador · Golden Kubestronaut



Bradley Andersen
Community Manager and Maintainer, k8gb
Kubestronaut

Kubernetes is expanding to the edge

Kubernetes on EDGE DAY NORTH AMERICA

Hundreds → thousands of locations. Not just cloud.

WHERE KUBERNETES RUNS



Retail



Telecom



Industrial sites

Edge reality: harsh, uncontrolled



Intermittent connectivity

Links to the outside world come and go.



Autonomy required

Process and decide at the source.



Strict latency

Users expect the nearest healthy site.

Every site is a failure domain. Traffic becomes the hard problem.

When a site goes dark, who steers the next request?

Kubernetes on EDGE DAY NORTH AMERICA

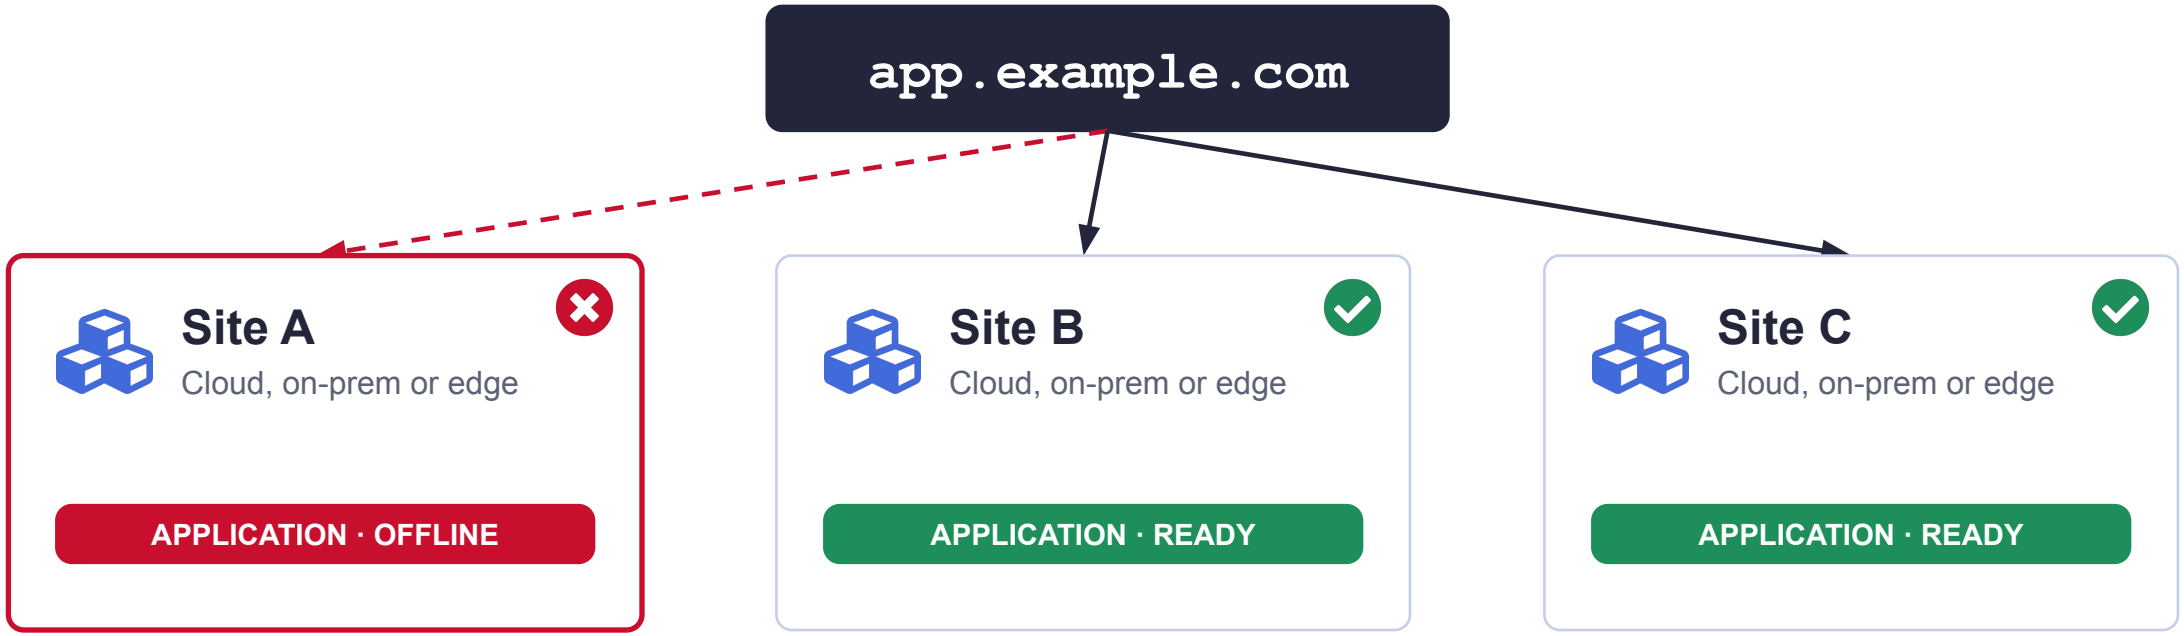
NODES
Spread replicas
Survive a worker failure

→

ZONES
Spread across zones
Survive a zone failure

→

CLUSTERS
Steer across clusters
Keep one hostname



Replicas and zones are not enough. Which healthy cluster gets the next request?

k8gb · Kubernetes Global Balancer

Kubernetes on EDGE DAY NORTH AMERICA



CNCF INCUBATING

Global application resilience across Kubernetes clusters. Two or more clusters, no fixed cluster-count limit.

Kubernetes-native

Readiness becomes DNS eligibility.

Decentralized

A control loop in each participating cluster. No central brain.

Declarative

Traffic steering is a custom resource you review in Git.

PRODUCTION ORIGINS
Born at Absa

30 MAR 2021
CNCF Sandbox

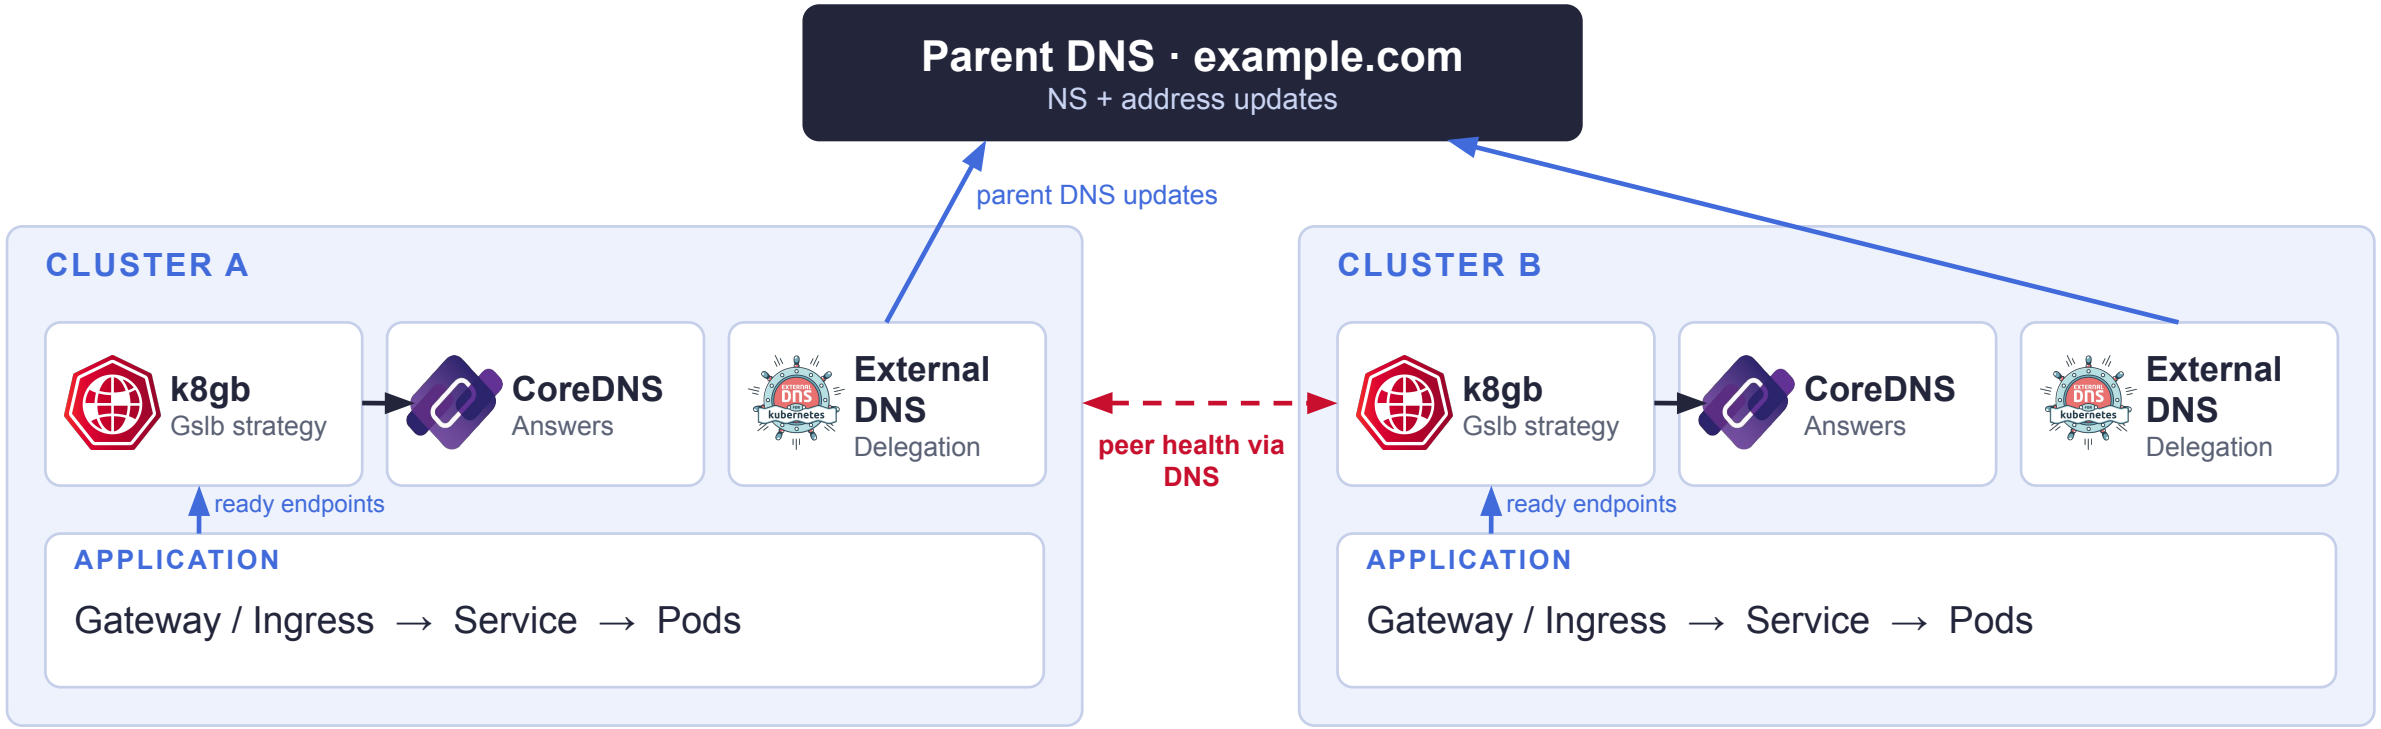
PUBLIC ADOPTION
Beyond Absa

18 JUL 2026
CNCF Incubating

Born from a production problem. Built with the community.

Independent clusters, shared DNS contract

Kubernetes on EDGE DAY NORTH AMERICA



k8gb controller Readiness + strategy → eligible targets

CoreDNS Delegated zone → authoritative answers

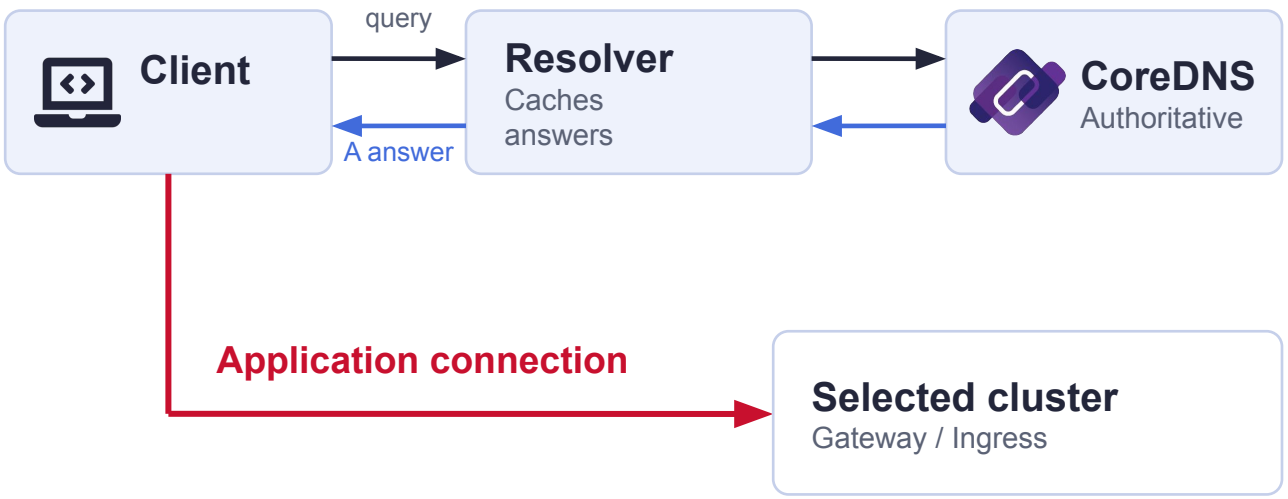
ExternalDNS Delegation state → parent DNS updates

Decide locally. Answer authoritatively. Delegate explicitly.

We don't proxy traffic: DNS selects, client connects

Kubernetes on EDGE DAY NORTH AMERICA

THE REQUEST PATH



DNS is outside the application data path.

No proxy hop and no extra latency on requests. Resolver caching means changes reach clients after TTL expiry.

DELEGATION IS THE AUTHORITY BOUNDARY

PARENT ZONE **example.com**

NS records for gslb.example.com → ns-a, ns-b (addresses published too)

DELEGATES

DELEGATED ZONE **gslb.example.com**

CoreDNS
Cluster A · ns-a

CoreDNS
Cluster B · ns-b

`A  app.gslb.example.com`
→ eligible application IPs

Own the name, the delegation, and DNS recovery.

Choose how eligible clusters are used

Kubernetes on EDGE DAY NORTH AMERICA

FAILOVER

Prefer primary



Return to primary on recovery.

ROUND ROBIN

Use all healthy clusters



Active / active DNS answers.

WEIGHTED ROUND ROBIN

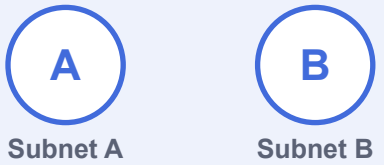
Express a preference



Request shares are approximate.

GEOIP

Map client subnets to clusters



Needs GeoIP data + ECS.

All modes change DNS answers, not existing connections.

Traffic steering is a declarative manifest



Kubernetes on EDGE DAY
NORTH AMERICA

gslb.yaml

```yaml
apiVersion: k8gb.io/v1beta1
kind: Gslb
metadata:
  name: test-gslb-failover
  namespace: test-gslb
spec:
  resourceRef:
    apiVersion: gateway.networking.k8s.io/v1
    kind: HTTPRoute
    name: payments
  strategy:
    type: failover
    primaryGeoTag: eu
```

Prefer A · eu

While it has eligible targets.

Same Gslb in every cluster

Each cluster observes health locally.

Gateway API first

HTTPRoute shown. Ingress, Istio VirtualService and LoadBalancer Service are also supported.

Prerequisites: the route, the zone, and matching geotags.

You define what a “site” means

Kubernetes on EDGE DAY NORTH AMERICA

Global traffic management can follow your own topology, down to a single edge box.

ILLUSTRATIVE clusterGeoTag VALUES

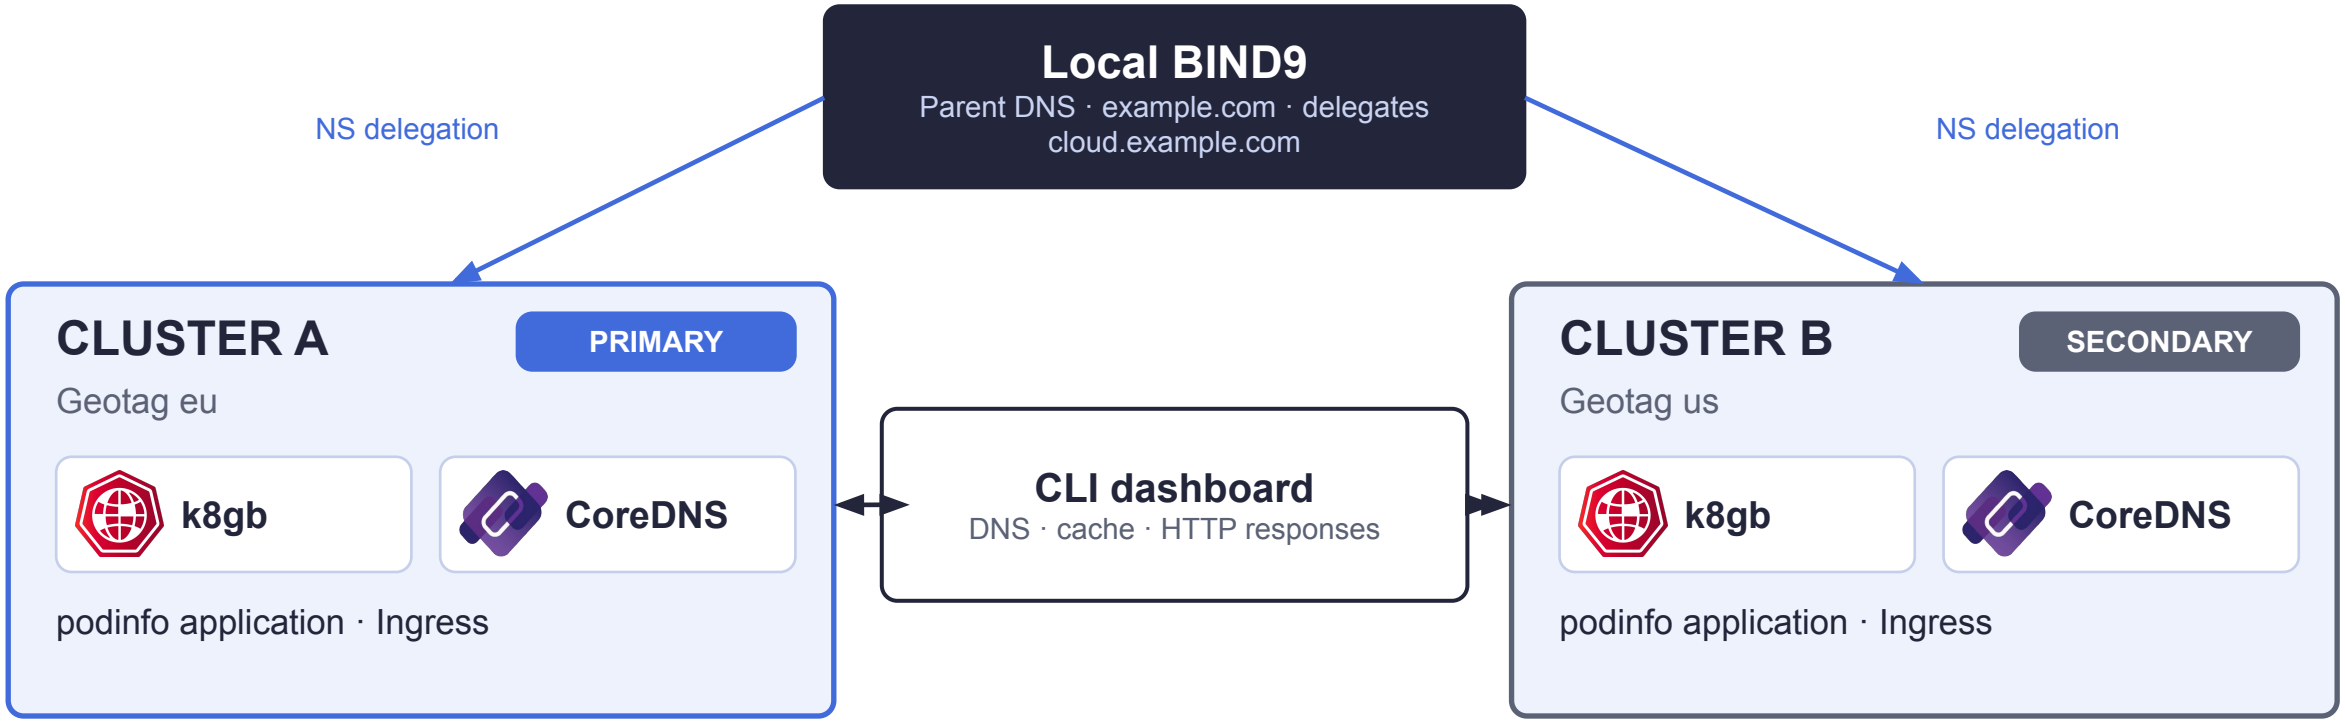
REGION	DATACENTER	RACK	SERVER	EDGE SITE
`eu-west`	`dc-1`	`rack-a`	`server-01`	`store-0412`

Each tag identifies a k8gb cluster.
Rack, server or edge-site boundaries use independent clusters, including single-node clusters.

Choose failure domains that match your edge topology.

Live demo · One hostname. Two clusters.

Kubernetes on EDGE DAY NORTH AMERICA



1 Observe A → 2 Remove A's ready targets → 3 Watch B take over → 4 Restore A

All local. Two application clusters represent two different sites.

A new DNS answer is not an instant switch

Kubernetes on EDGE DAY NORTH AMERICA

Detection, DNS publication and client recovery happen at different times.



TIMING VARIES

Readiness · reconciliation · resolver caches · client retries

New DNS answer ≠ recovered clients

Measure recovery at the resolver and at the application.

Readiness must reflect the real operation



Kubernetes on EDGE DAY
NORTH AMERICA

`readinessProbe`
App-defined check

→

Pod Ready
Kubernetes condition

→

EndpointSlice
Service endpoint readiness

→

Gslb targets
Readiness + strategy

HEALTH ENDPOINT

`GET /healthz`

HTTP 200

The shallow check passes.

REAL REQUEST

Payment cannot be saved

`POST /payments`
→ **database is READ-ONLY**

Replication and safe writer ownership belong to the application.

k8gb · traffic steering
Excludes unready targets; publishes healthy clusters.

Yours · stateful data & sessions
Replication, fencing, recovery and session continuity.

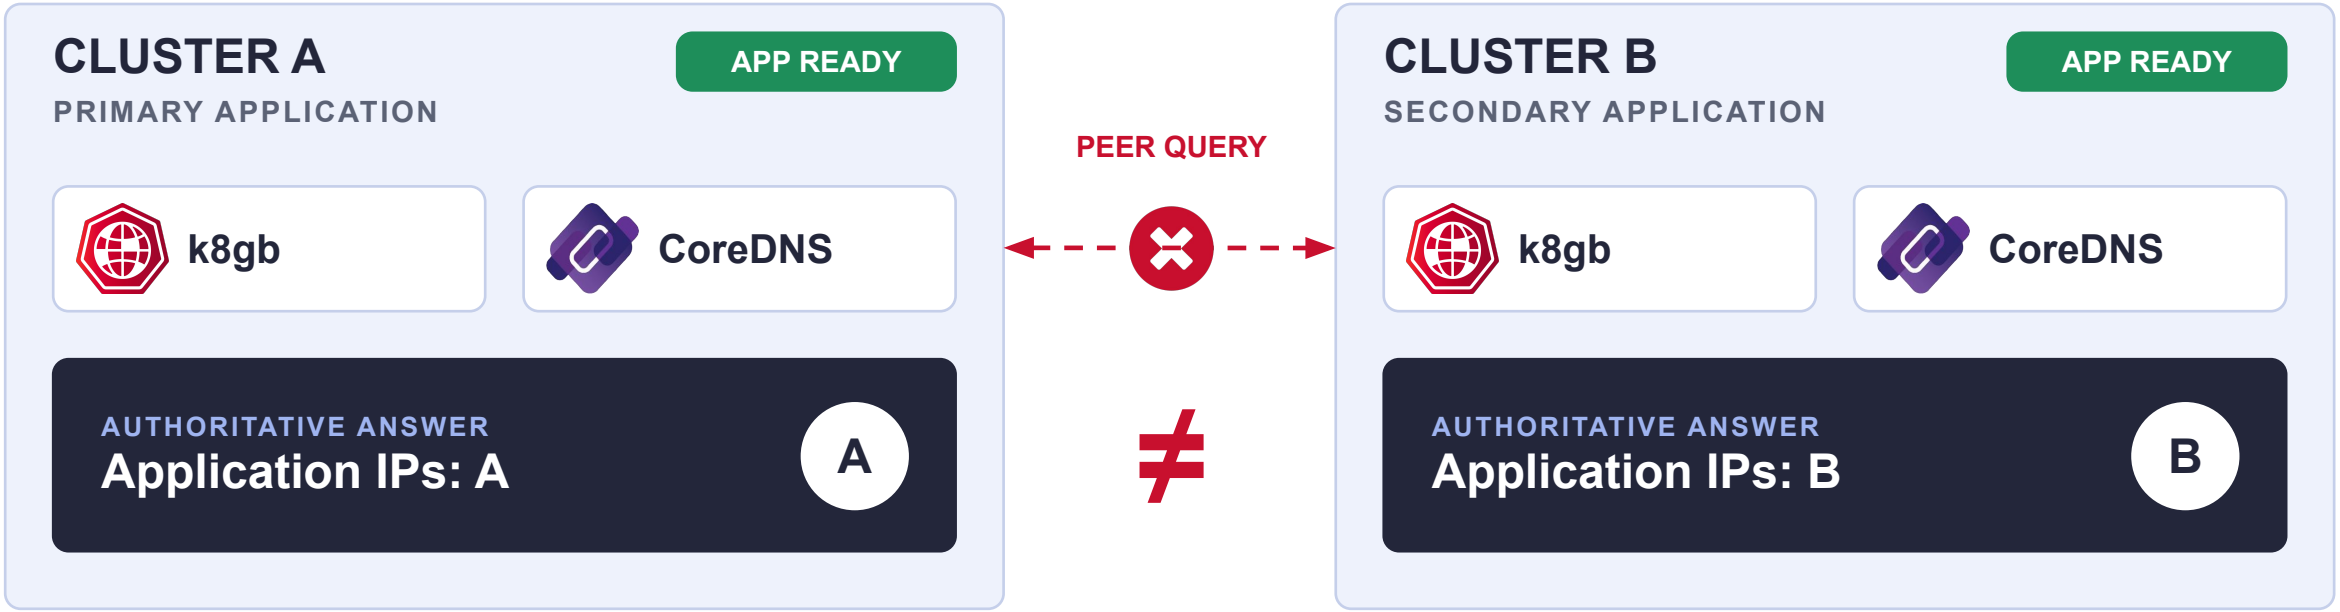
Traffic recovery ≠ data recovery ≠ session continuity.

k8gb steers traffic. The application must be ready to serve.

At the edge, partitions are normal

Kubernetes on EDGE DAY NORTH AMERICA

Both applications are ready. Clients can still reach both DNS authorities. The clusters cannot reach each other.



B cannot obtain A's peer-health answer, so the two authorities publish different answers.

A failed health query does not prove the application is down.

Why this fits the edge, and what to test

Kubernetes on EDGE DAY NORTH AMERICA

Edge need	What k8gb provides	Watch out for
Autonomy	A control loop + authoritative CoreDNS in every cluster	Parent-zone delegation is a shared dependency
Handling failures	Unready endpoints drop out of DNS answers	Readiness must reflect the real operation
Low latency	GeoIP strategy maps client subnets to clusters	Needs GeoIP data + ECS
Intermittent connectivity	Each authority answers from what it can observe; agreement after reconciliation	During a partition, answers can differ
Your topology	clusterGeoTag can be a region, DC, rack, server or site	One tag = one independent cluster

Test your traffic strategy in your environment

Define expectations
Expected routing for each scenario.

Rehearse scenarios
Failures and recovery, including partitions.

Verify the result
DNS answers, real requests and timing.

Automate and repeat
Run the checks after every change.

Validate failover, failback, and every Gslb strategy you rely on.

In production, built by a community

Kubernetes on EDGE DAY NORTH AMERICA

MILLENNIUM BCP · CNCF CASE STUDY

70%

shorter response times for DNS-related incidents

5 regions
Two cloud providers + on-premises

Zero downtime
Reported during regional failover tests

A multi-region deployment, not an edge one: the same mechanism.

Source: CNCF case study, 16 September 2025 · cncf.io/case-studies/millennium-bcp

7	46	1,335
public adopters	contributors	GitHub stars

CNCF / GitHub, checked 23 October 2026

PUBLIC ADOPTERS

Absa · Millennium bcp · Eficode · Open Systems · PagBank · Darede · envio.dev

CNCF Incubating since July 2026

Meet the community every other Wednesday. Everyone is welcome.

Toward a sovereign reference stack

Kubernetes on EDGE DAY NORTH AMERICA

NeoNephos reference architecture NeoNephos (Linux Foundation Europe) hosts open-source building blocks for a sovereign European cloud-to-edge continuum, including IronCore for Kubernetes-style bare-metal management.

01

Prove viability

Run k8gb on IronCore in a Box, locally.

02

Real bare metal

Deploy the k8gb AI inference demo on Hetzner bare metal, with k8gb.de and k8gb.fi as targets.

03

Share it

A blog post and a KubeCon talk.

HELP WANTED

Pick it up

k8gb issue #2535

Help us build a sovereign k8gb reference architecture.

Questions?

Kubernetes on EDGE DAY NORTH AMERICA

Try k8gb, meet the maintainers, share your deployment.

Star & contribute

github.com/k8gb-io/k8gb

Read the docs

k8gb.io

Join the conversation
#k8gb on CNCF Slack

Community meetings
Every other Wednesday, everyone welcome

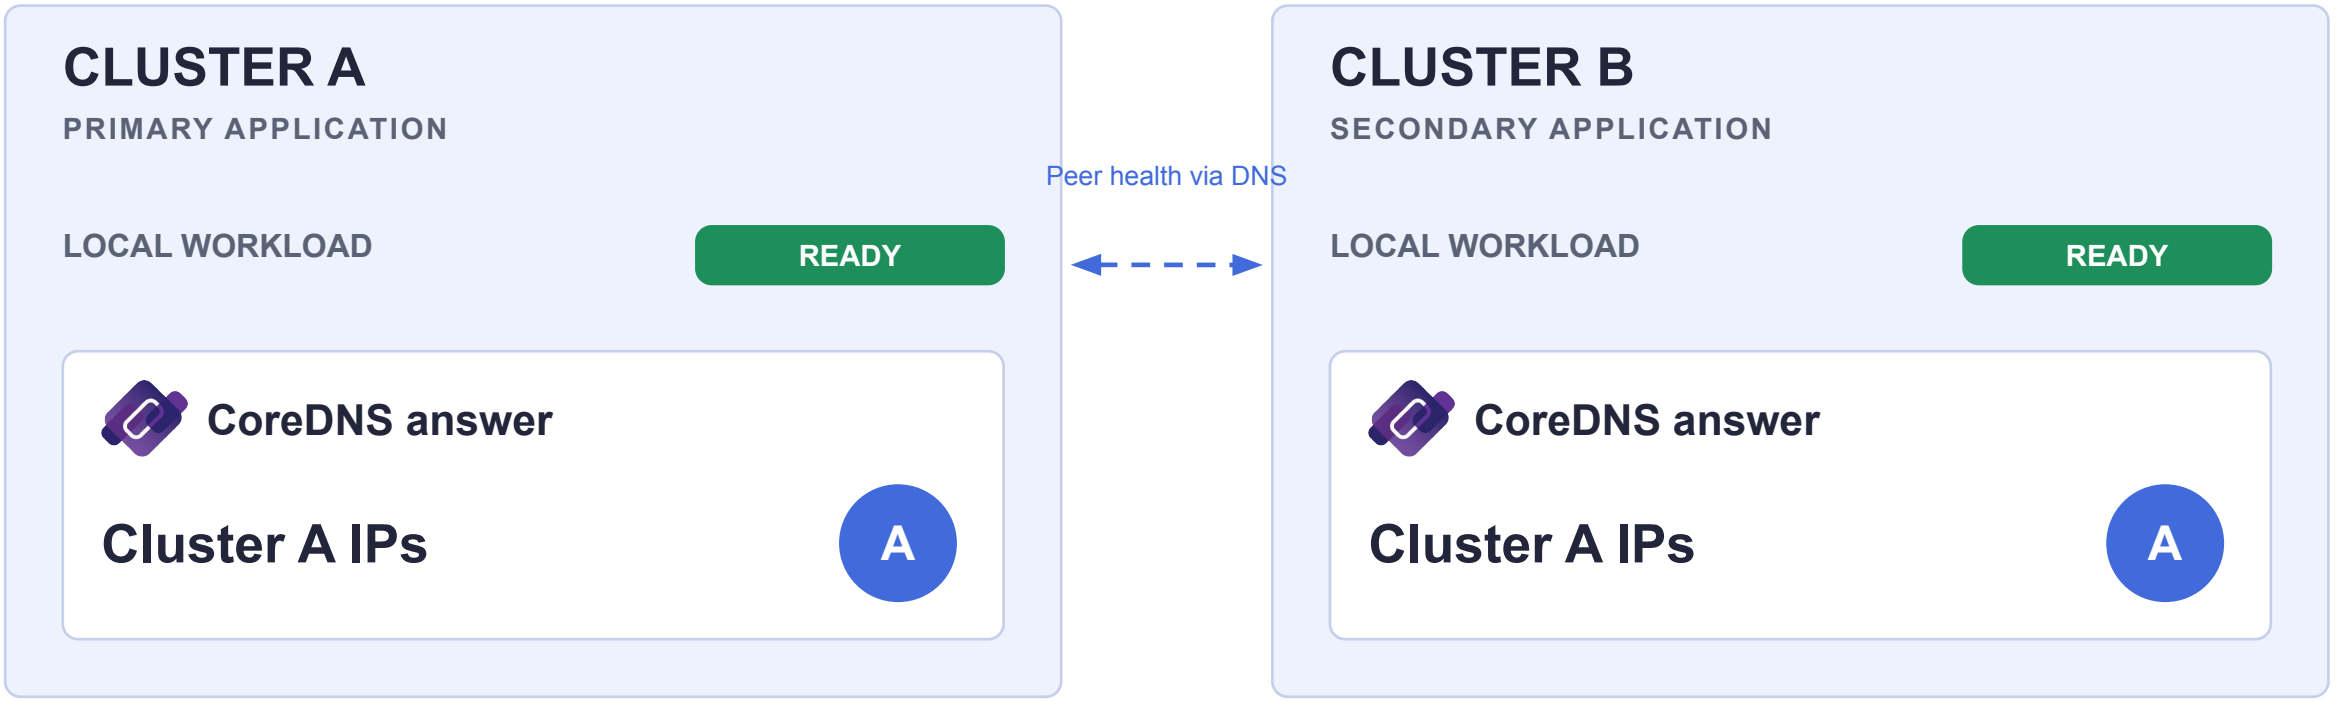
Share your deployment
Add your organization to ADOPTERS.md

Fabrizio Sgura · fabrizio-sgura.bsky.social | Bradley Andersen · bradley.andersen@pm.me

Backup · Demo recap 1: normal operation

Kubernetes on EDGE DAY NORTH AMERICA

Both applications are ready. Cluster A is the preferred target.



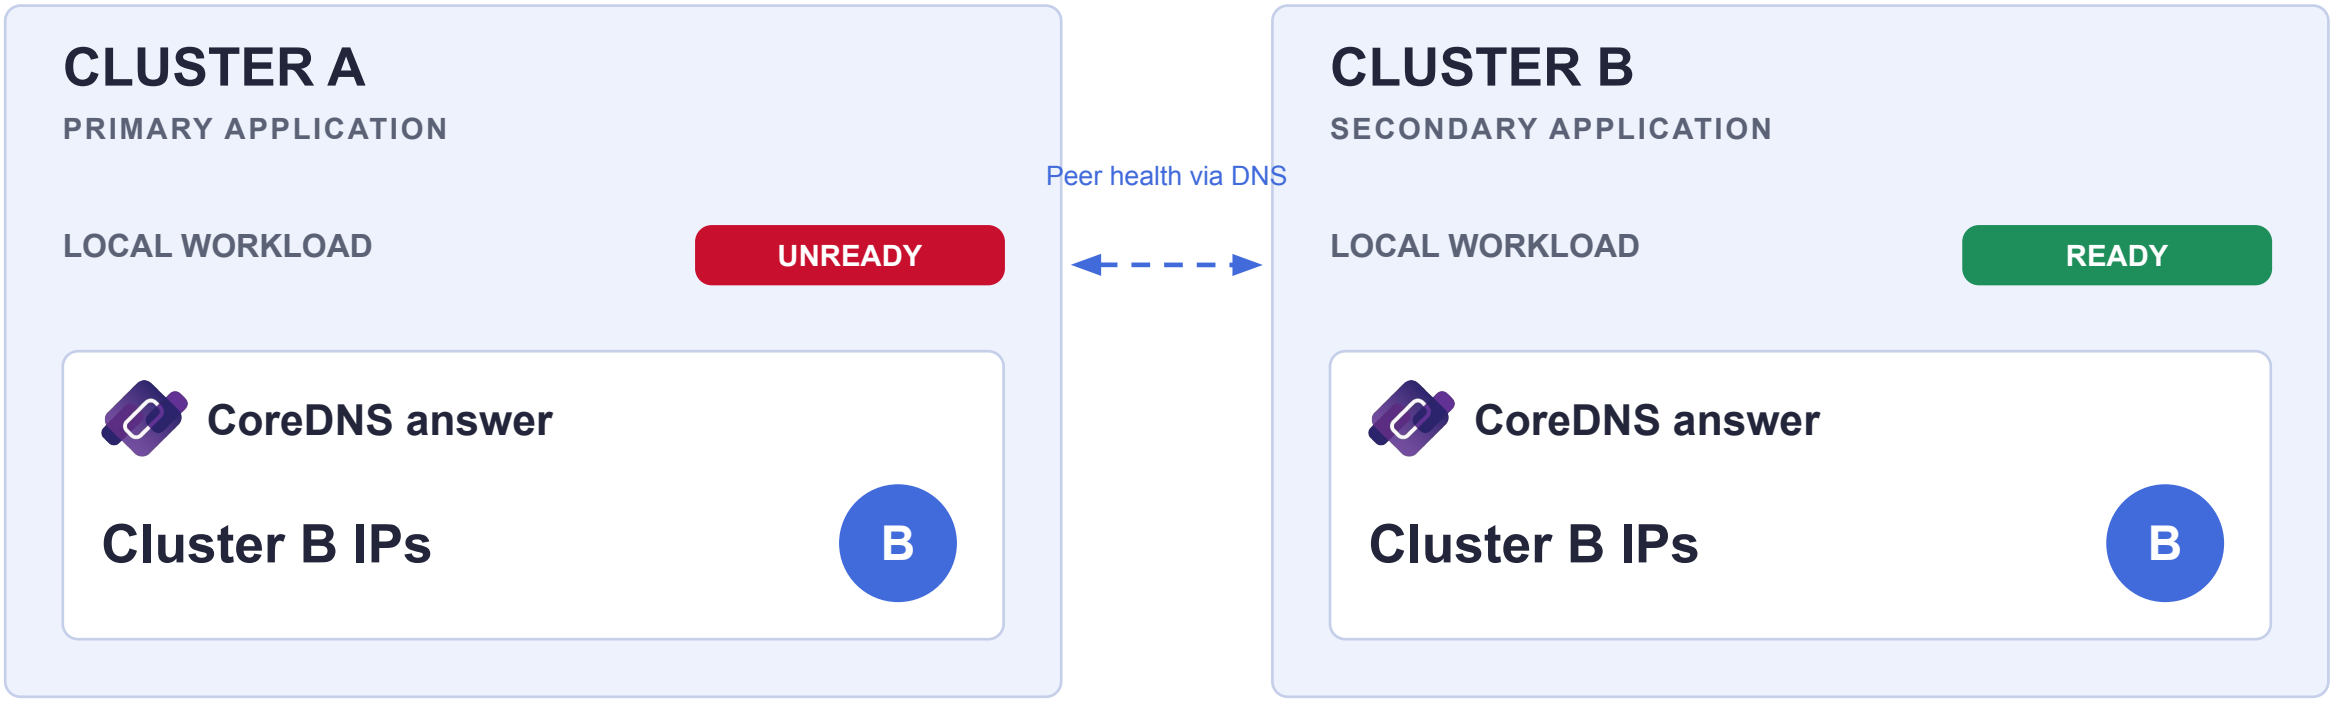
Each k8gb controller queries peer CoreDNS.

Both authorities agree after reconciliation; client caches may lag.

Backup · Demo recap 2: A loses ready targets

Kubernetes on EDGE DAY NORTH AMERICA

A's application is unready; its DNS authority is still available. Both publish B.



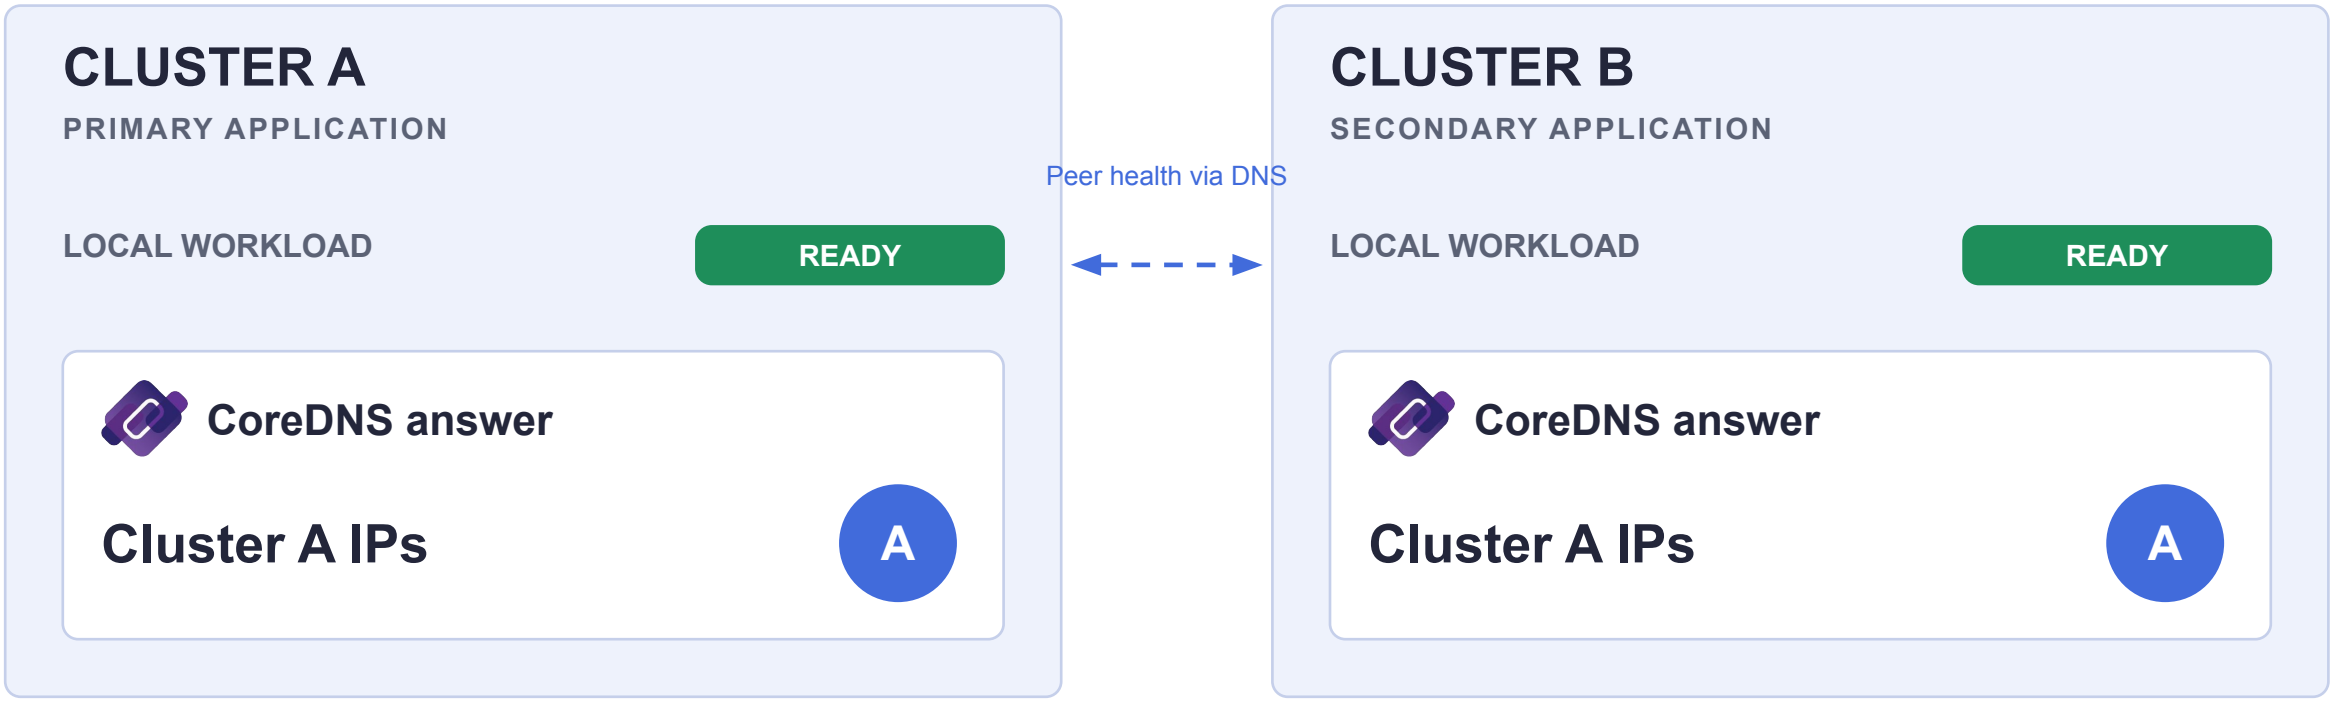
Each k8gb controller queries peer CoreDNS.

Both authorities agree after reconciliation; client caches may lag.

Backup · Demo recap 3: A recovers

Kubernetes on EDGE DAY NORTH AMERICA

A's application is ready again. The primary preference returns.



Each k8gb controller queries peer CoreDNS.

Both authorities agree after reconciliation; client caches may lag.

Backup · Kubernetes and DNS integrations



Kubernetes on EDGE DAY
NORTH AMERICA

Run on-prem, in the cloud, or across both.

NATIVE RESOURCE INTEGRATIONS

- Gateway API
- Ingress
- Istio VirtualService
- LoadBalancer Service

on any conformant Kubernetes

AUTOMATED PARENT-ZONE DELEGATION

AWS Route 53	**Azure Public DNS**	**Google Cloud DNS**
NS1	**Infoblox**	**Cloudflare**

RFC 2136 DNS updates
BIND9 · Windows DNS · TSIG / GSS-TSIG authentication

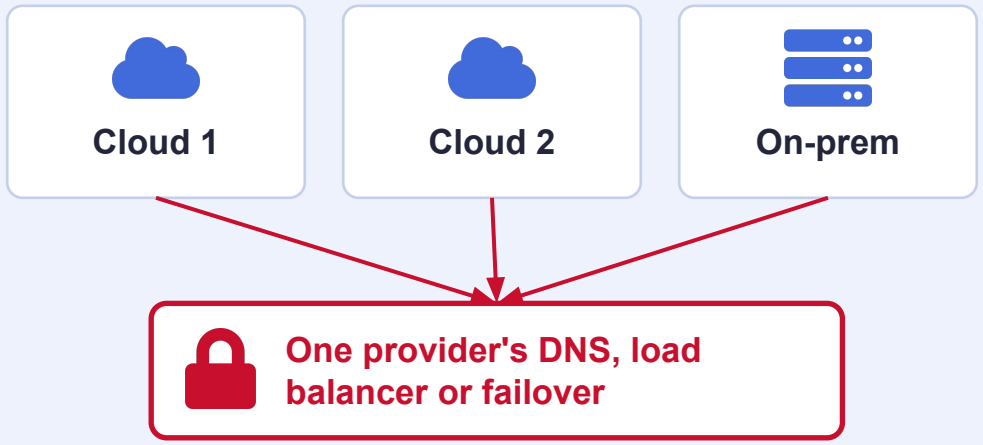
Choose your parent DNS. Keep your Gslb traffic strategy.

Backup · Sovereignty: who controls the failover?

Kubernetes on EDGE DAY NORTH AMERICA

Sovereignty is also about who controls reachability when infrastructure fails, not only where workloads and data live.

THE HIDDEN DEPENDENCY



Distributed across providers, yet failover still depends on one of them.

K8GB: OPEN, PROVIDER-INDEPENDENT FAILOVER

Environment agnostic

Conformant Kubernetes: cloud, on-premises or hybrid.

Distributed DNS

k8gb + authoritative CoreDNS in every participating cluster.

Own your DNS

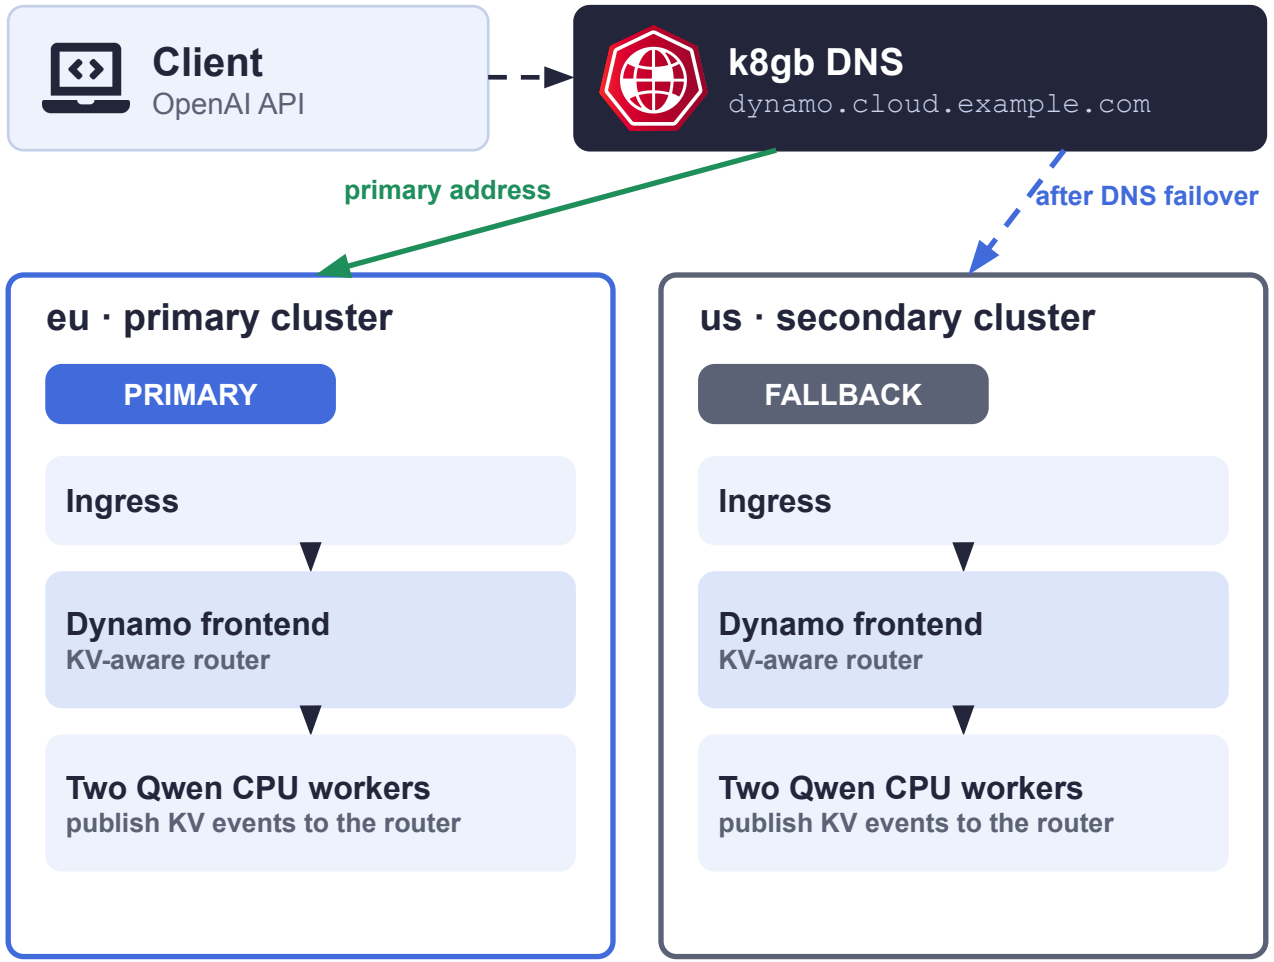
Self-hosted BIND9 (RFC 2136) for the parent zone and delegation.

Is your platform sovereign, or only distributed across several providers?

Your infrastructure. Your DNS. Your failover.

Backup · AI inference with NVIDIA Dynamo

Kubernetes on EDGE DAY NORTH AMERICA



WHO DOES WHAT

Select a healthy regional entry	**k8gb Gslb failover**
Choose a local worker per request	**Dynamo KV router**
Withdraw an unavailable model	**Kubernetes readiness probe**
Reconcile frontend + workers	**Dynamo operator (DGD)**

DRAFT PR #2579

Dynamo 1.6.0 + vLLM, Qwen3-0.6B on CPU, no GPU.

Validated locally: failover to us and failback to eu on one hostname; repeated prompts routed to warm workers.

k8gb selects a healthy region through DNS; Dynamo selects a worker using local cache state.