

Google Cloud

Next '24

Introducing Cloud
Service Mesh: A
fully managed
global scale
service mesh



**Ameer
Abbas**

Product Manager,
Google Cloud



**Chris
Crall**

Product Manager,
Google Cloud



**Arunkumar
Jayaraman**

Technical Lead
Service Mesh,
Google Cloud

Agenda

- 01 Why service mesh?
- 02 Cloud Service Mesh
- 03 Evolution of mesh
- 04 Ambient Mesh
- 05 API
- 06 Extensibility
- 07 Istio
- 08 Demo

Why Service Mesh?

Scale and Reliability

Run **globally scalable**
and reliable applications



Why Service Mesh?

Scale and Reliability

Run **globally scalable and reliable** applications

Security and Policy

Architect **zero trust and policy driven** networks



Why Service Mesh?

Scale and Reliability

Run **globally scalable and reliable** applications

Security and Policy

Architect **zero trust and policy driven** networks

Service Management

Service-centric telemetry and application management

Complexity and supportability of Mesh

Complexity

Onboarding

Lifecycle Management

Integration with other Services

Supportability

Self supported

Strong Community Reliance

Large scale production

support challenges

Reliability

Requires strong community

Requires strong adoption

Consistent innovation and

features

Google Service Mesh products

Anthos Service Mesh

Istio API based (Kubernetes CRDs)

Google Cloud and non Google Cloud environments

Managed and *hosted* on GCP Managed and *local* on non GCP

Platform Admin/Service Operators

Traffic Director

GCP API based

Google Cloud environments

Multiple GCP compute runtime support

GCP networking services integration

Managed and hosted on GCP

Network Admins



Introducing Cloud Service Mesh

A globally scalable, fully managed, Google platform integrated service mesh for all enterprises



Global control plane
Run anywhere

A globally scalable, fully managed, Google platform integrated service mesh for all enterprises

Global control plane
Run anywhere

Google managed
control/data plane &
Mesh/CA Services

A globally scalable, fully
managed, Google platform
integrated service mesh for
all enterprises

Global control plane
Run anywhere

Google managed
control/data plane &
Mesh/CA Services

A globally scalable, fully managed, Google platform integrated service mesh for all enterprises

Integrated with
Google
Networking
Services

Global control plane
Run anywhere

Google managed
control/data plane &
Mesh/CA Services

A globally scalable, fully managed, Google platform integrated service mesh for all enterprises

Integrated with
Google
Networking
Services and
compute
platforms

Enterprise grade mesh to
run secure, reliable and
managed services

**Managed
Runtimes**

GKE/E

GCE

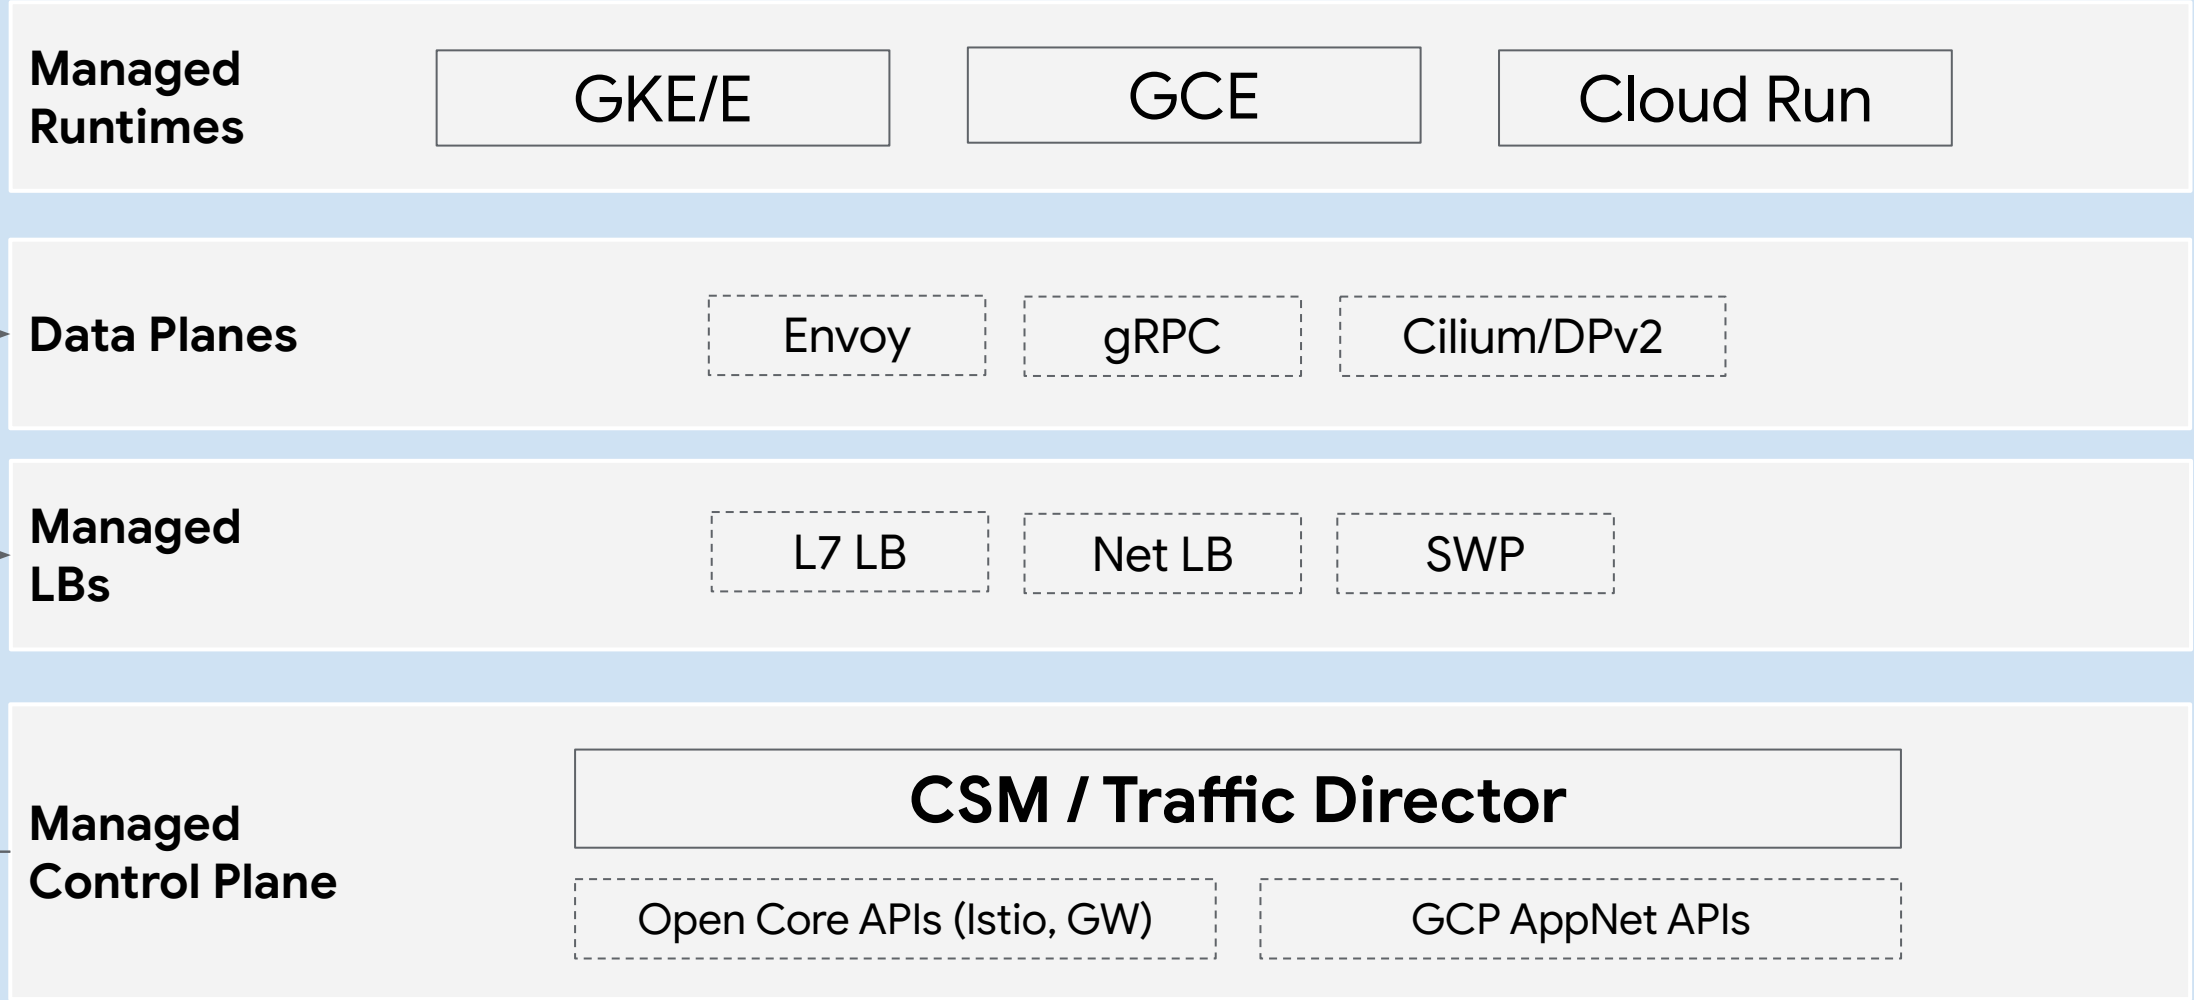
Cloud Run

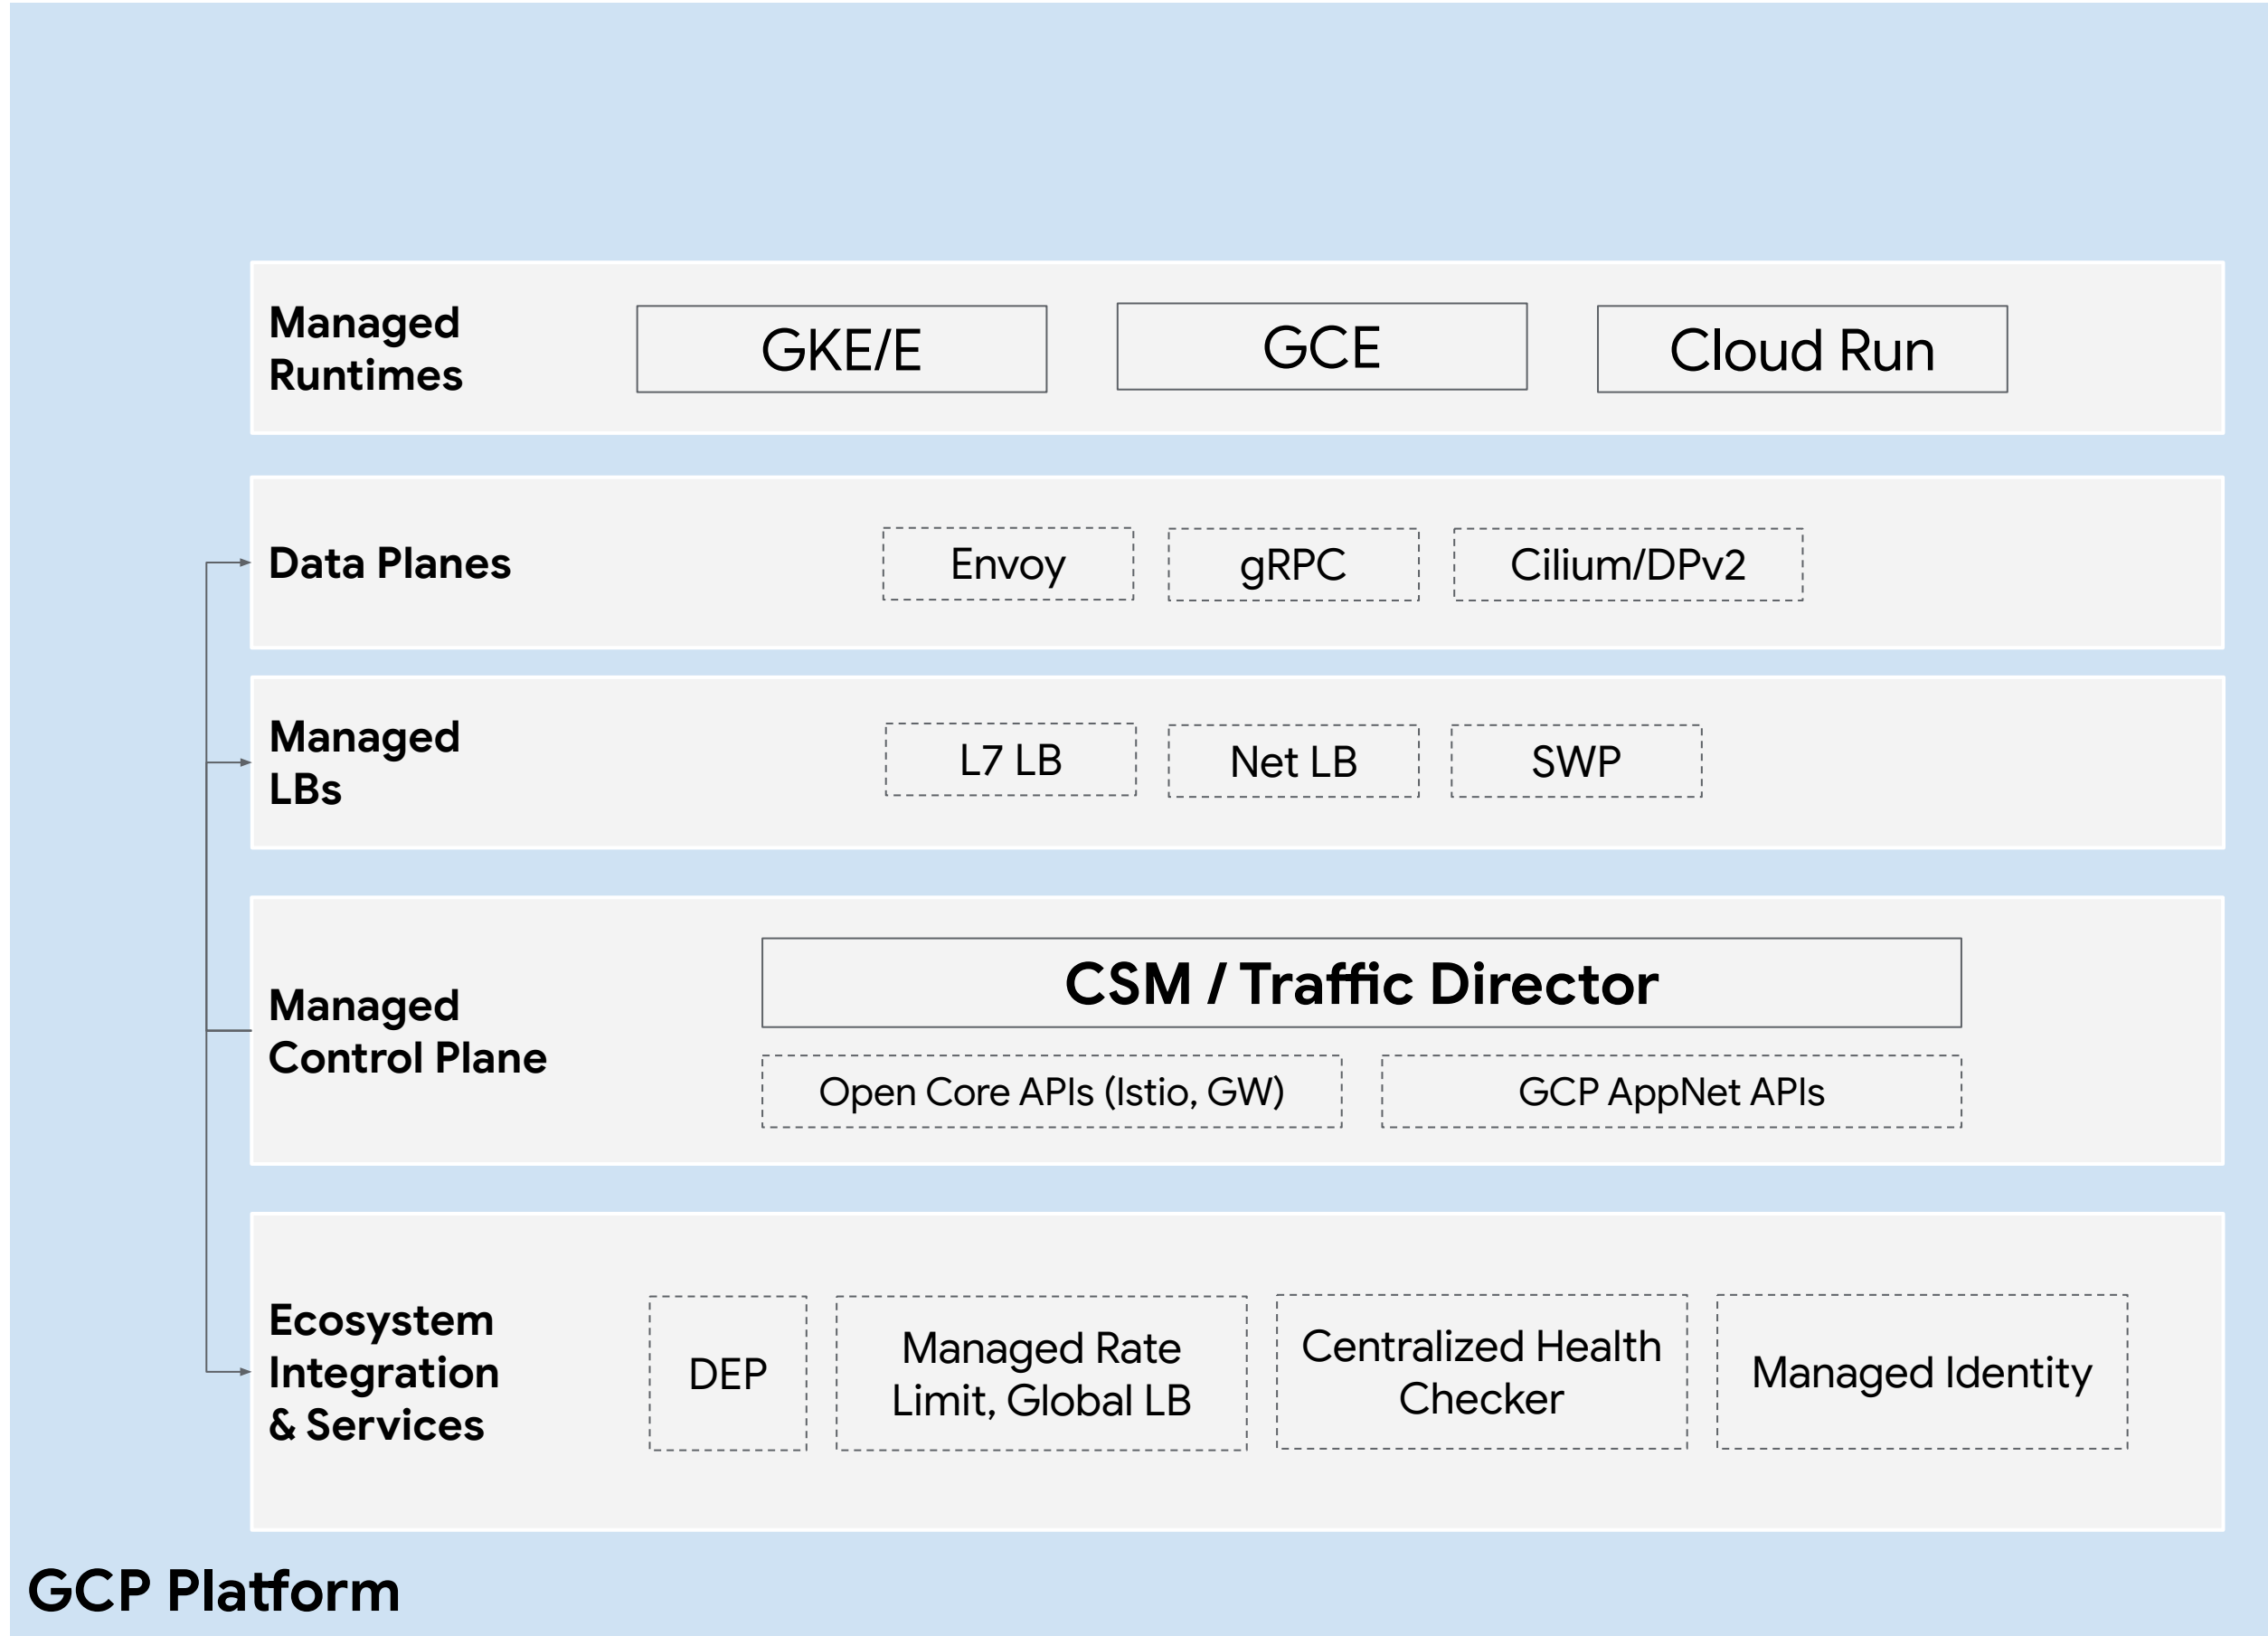
**Managed
LBs**

L7 LB

Net LB

SWP





Customer
Manages

Apps, Services, Workloads

**Application
Management**

App Hub

**Managed
Runtimes**

GKE/E

GCE

Cloud Run

Data Planes

Envoy

gRPC

Cilium/DPv2

**Managed
LBs**

L7 LB

Net LB

SWP

**Managed
Control Plane**

CSM / Traffic Director

Open Core APIs (Istio, GW)

GCP AppNet APIs

**Ecosystem
Integration
& Services**

DEP

Managed Rate
Limit, Global LB

Centralized Health
Checker

Managed Identity

GCP Platform

Customer
Manages

Apps, Services, Workloads

**Application
Management**

App Hub

Services, Workloads

**Managed
Runtimes**

GKE/E

GCE

Cloud Run

On-prem

Multi-cloud

Data Planes

Envoy

gRPC

Cilium/DPv2

Envoy

gRPC

**Managed
LBs**

L7 LB

Net LB

SWP

**Managed
Control Plane**

CSM / Traffic Director

Istiod

Open Core APIs (Istio, GW)

GCP AppNet APIs

Open Core APIs (Istio, GW)

**Ecosystem
Integration
& Services**

DEP

Managed Rate
Limit, Global LB

Centralized Health
Checker

Managed Identity

GCP Platform

Off-GCP

Customer
Manages

Apps, Services, Workloads

**Application
Management**

App Hub

Services, Workloads

**Managed
Runtimes**

GKE/E

GCE

Cloud Run

On-prem

Multi-cloud

Data Planes

Envoy

gRPC

Cilium/DPv2

Envoy

gRPC

**Managed
LBs**

L7 LB

Net LB

SWP

**Managed
Control Plane**

CSM / Traffic Director

Istiod

Open Core APIs (Istio, GW)

GCP AppNet APIs

Open Core APIs (Istio, GW)

**Ecosystem
Integration
& Services**

DEP

Managed Rate
Limit, Global LB

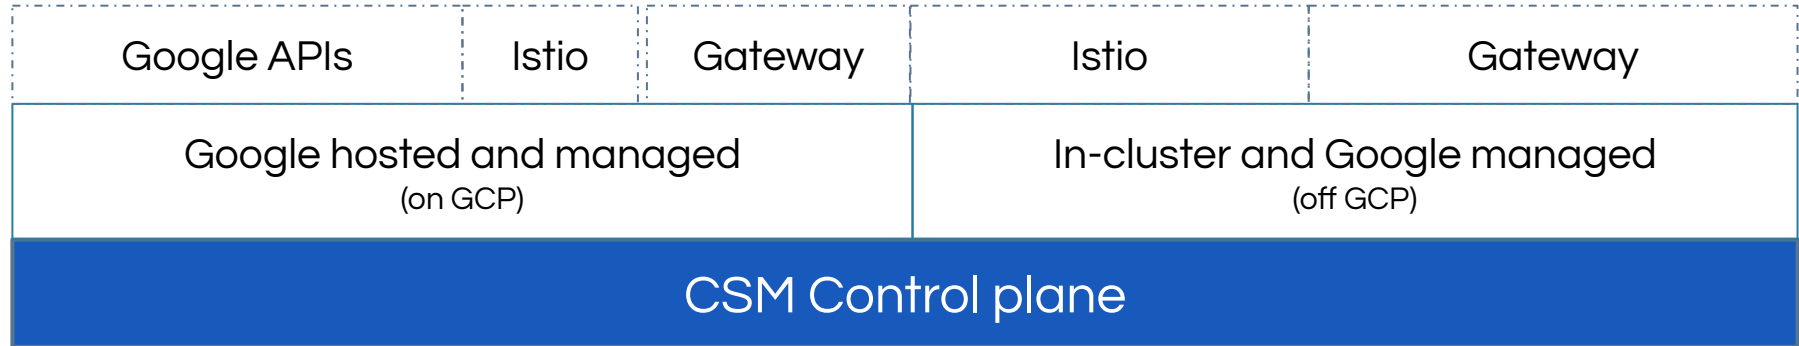
Centralized Health
Checker

Managed Identity

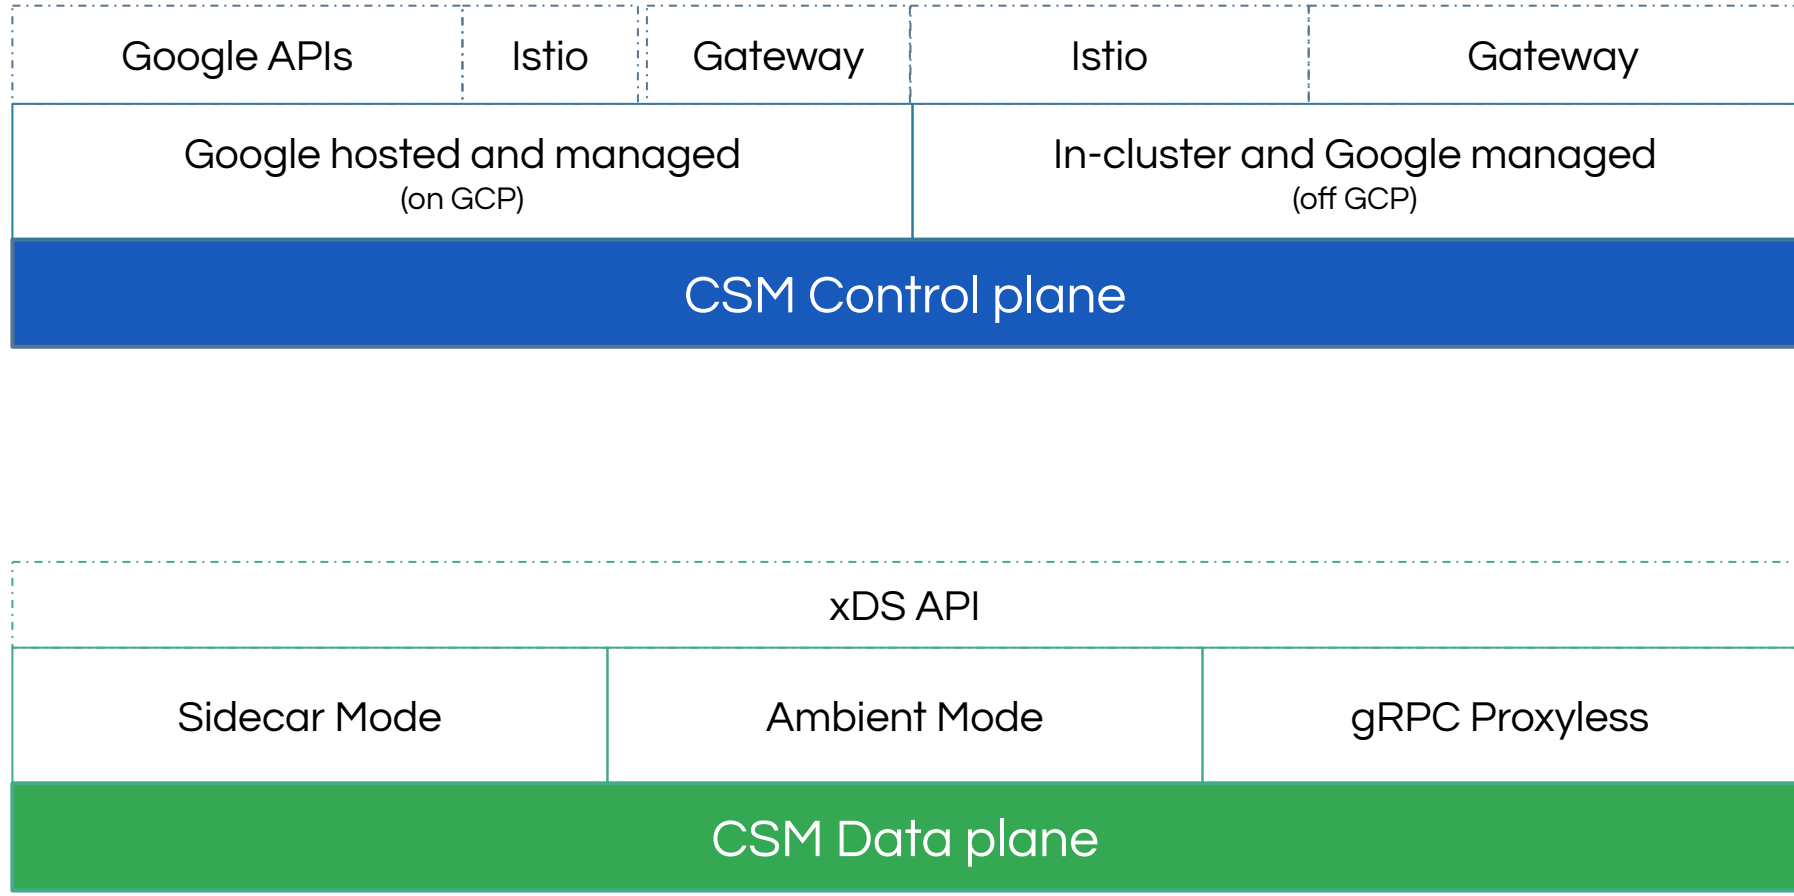
GCP Platform

Off-GCP



Cloud Service Mesh Architecture





Cloud Service Mesh Architecture



Cloud Service Mesh Architecture

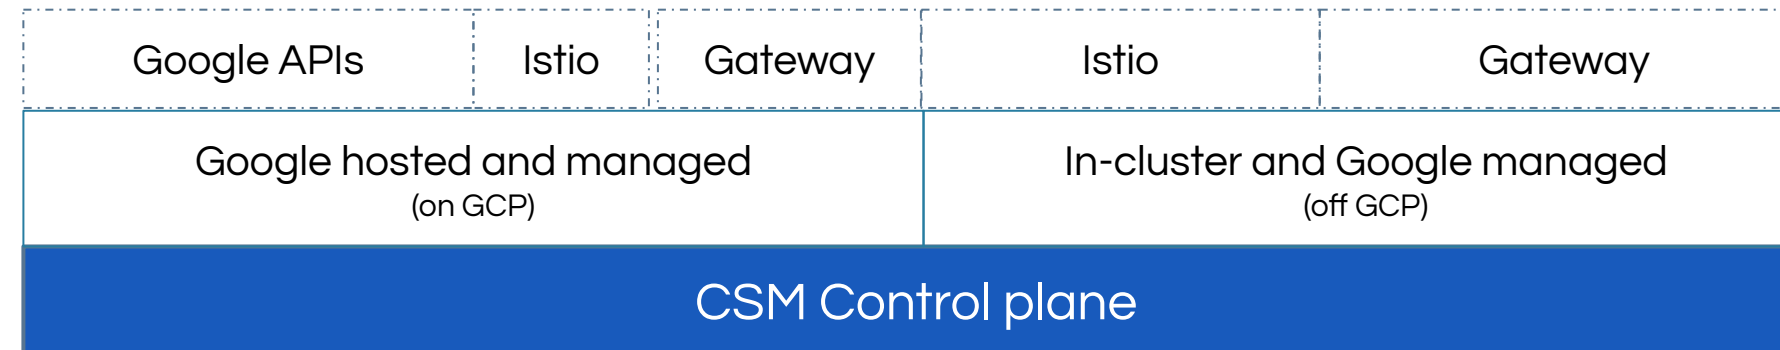


Cloud Service Mesh Architecture



Terraform
gcloud (API)
UI

CSM Control plane



On GCP

Hosted

Managed upgrades and patches

Supports Google, Istio and Gateway APIs

GCP platform integrations


Off GCP

In cluster

Managed upgrades and patches

Supports Istio, Gateway APIs

CSM Data plane



Sidecar Mode

Sidecar Proxy per workload

Proxy lifecycle is application disruptive

Higher resource consumption and latency

No shared proxies

Single proxy for L4 and L7 functions

Ambient Mode

Shared ultra light proxy per node

Shared proxy handles L4 (standard) functions

Optional per-service proxy handles L7 (advanced) functions

Non-disruptive to applications

Lower resource consumption

Better performance

gRPC Proxyless

xDS protocol in gRPC libraries

Integrated with gRPC application

Lighter weight (vs sidecar)

Better performance (vs sidecar)

mTLS to the App

CSM Mesh Services



Certificate Authority

Google-managed CA Service

- Option of two CAs

Self-managed in-cluster CA

Service and API Management

Services Dashboards

SLOs and Alerting

Cloud Operations (logs, tracing, custom dashboards)

Policy Controller

Security Insights

Anthos and Traffic Director Users

ASM for GCP

Managed Control Plane migrated to TD control plane (gradual, managed, reliable process)

No change to:

- Istio APIs
- Managed Data Plane
- Metrics
- CAs
- Service UI

Anthos and Traffic Director Users

ASM for GCP

Managed Control Plane migrated to TD control plane (gradual, managed, reliable process)

No change to:

- Istio APIs
- Managed Data Plane
- Metrics
- CAs
- Service UI

ASM Off-GCP

No changes short term

Direction:

- Remain Open Core APIs
- More managed environment => remove toil
- Federation with GCP control plane

Anthos and Traffic Director Users

ASM for GCP

Managed Control Plane migrated to TD control plane (gradual, managed, reliable process)

No change to:

- Istio APIs
- Managed Data Plane
- Metrics
- CAs
- Service UI

ASM Off-GCP

No changes short term

Direction:

- Remain Open Core APIs
- More managed environment => remove toil
- Federation with GCP control plane

Traffic Director

Existing customers see no changes.

Direction:

- Tighter integration with GCP Networking
- GCP API

Evolution of Mesh

Evolution of Mesh

GCP Integrations

GKE, Cloud Run, GCE VM

Global load balancing,
Centralized health checking

Traffic Driven Autoscaling

Global rate limiting

PSC

Evolution of Mesh

GCP Integrations

GKE, Cloud Run, GCE VM

Global load balancing,
Centralized health checking

Traffic Driven Autoscaling

Global rate limiting

PSC

API

GCP APIs

Open Core APIs

Evolution of Mesh

GCP Integrations

GKE, Cloud Run, GCE VM

Global load balancing,
Centralized health checking

Traffic Driven Autoscaling

Global rate limiting

PSC

API

GCP APIs

Open Core APIs

Multi Data Plane Support

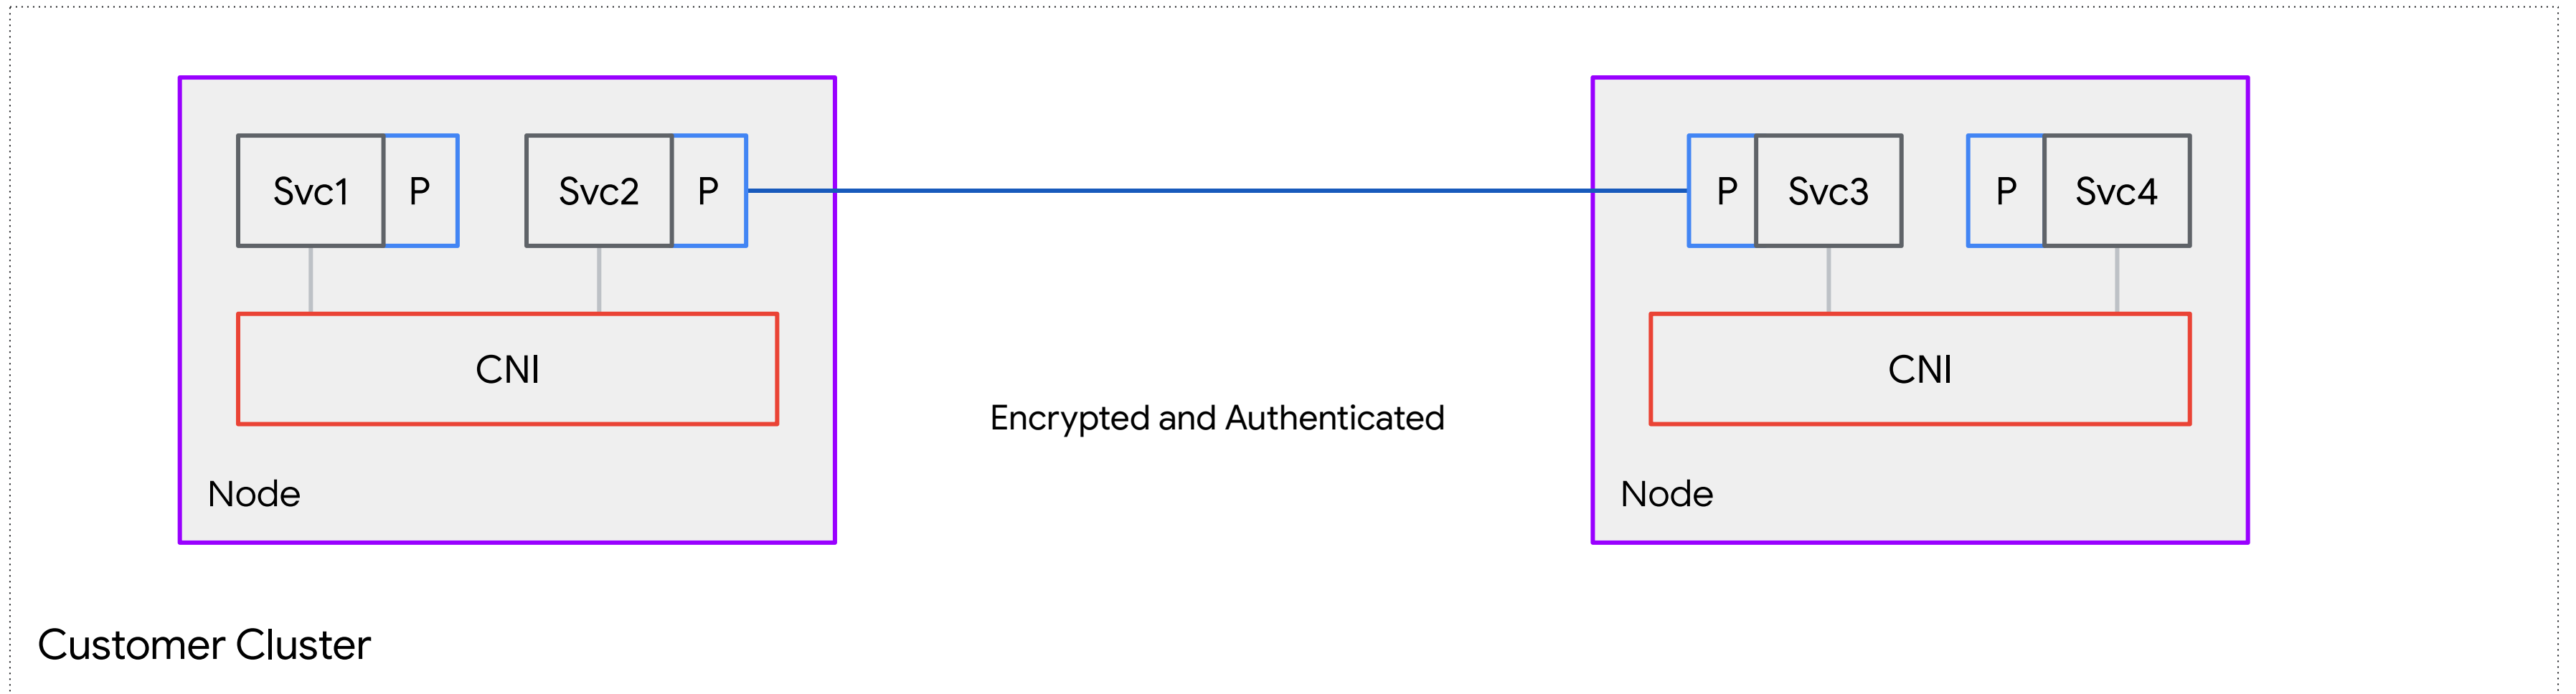
Envoy Sidecar

Proxyless gRPC

GKE native, Managed load balancer
integrated Ambient mode

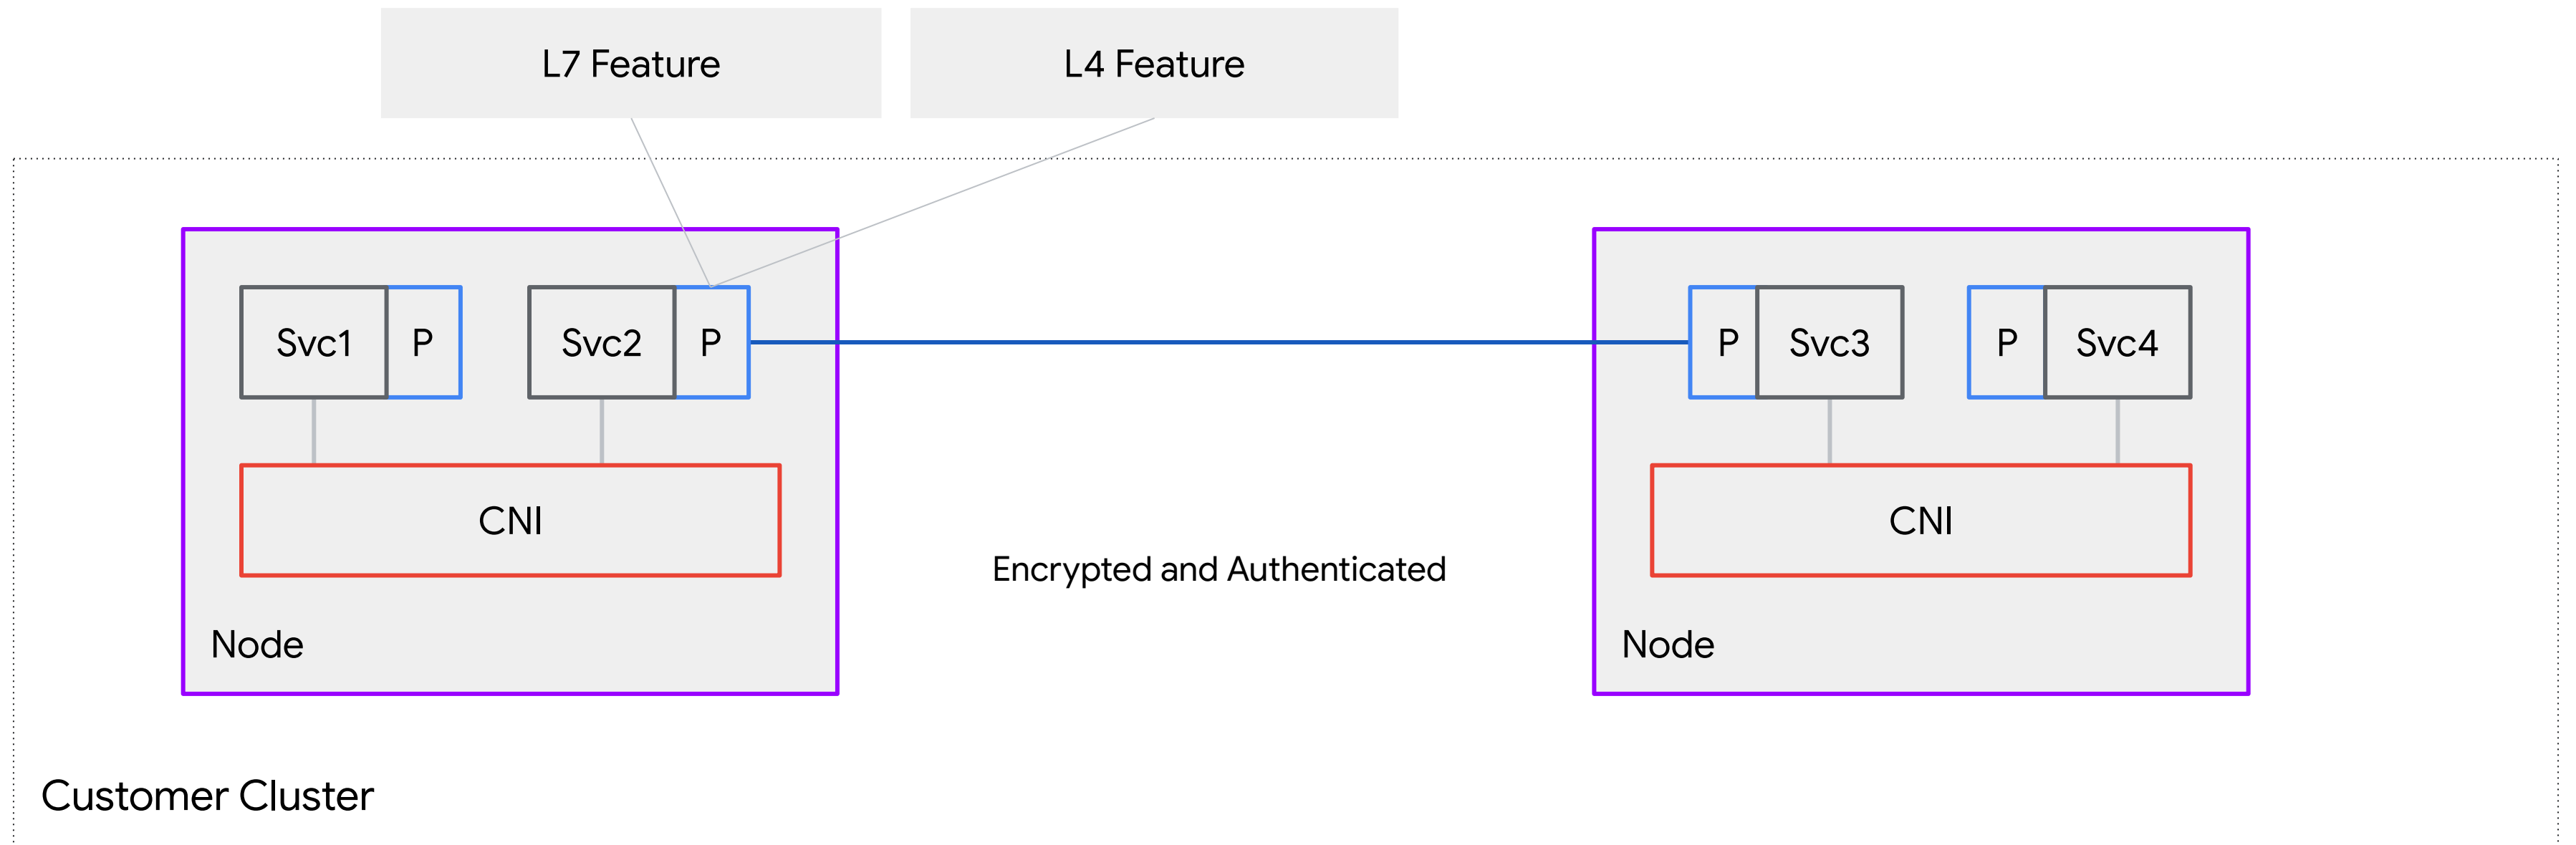
Ambient Mesh Architecture

Mesh with sidecars

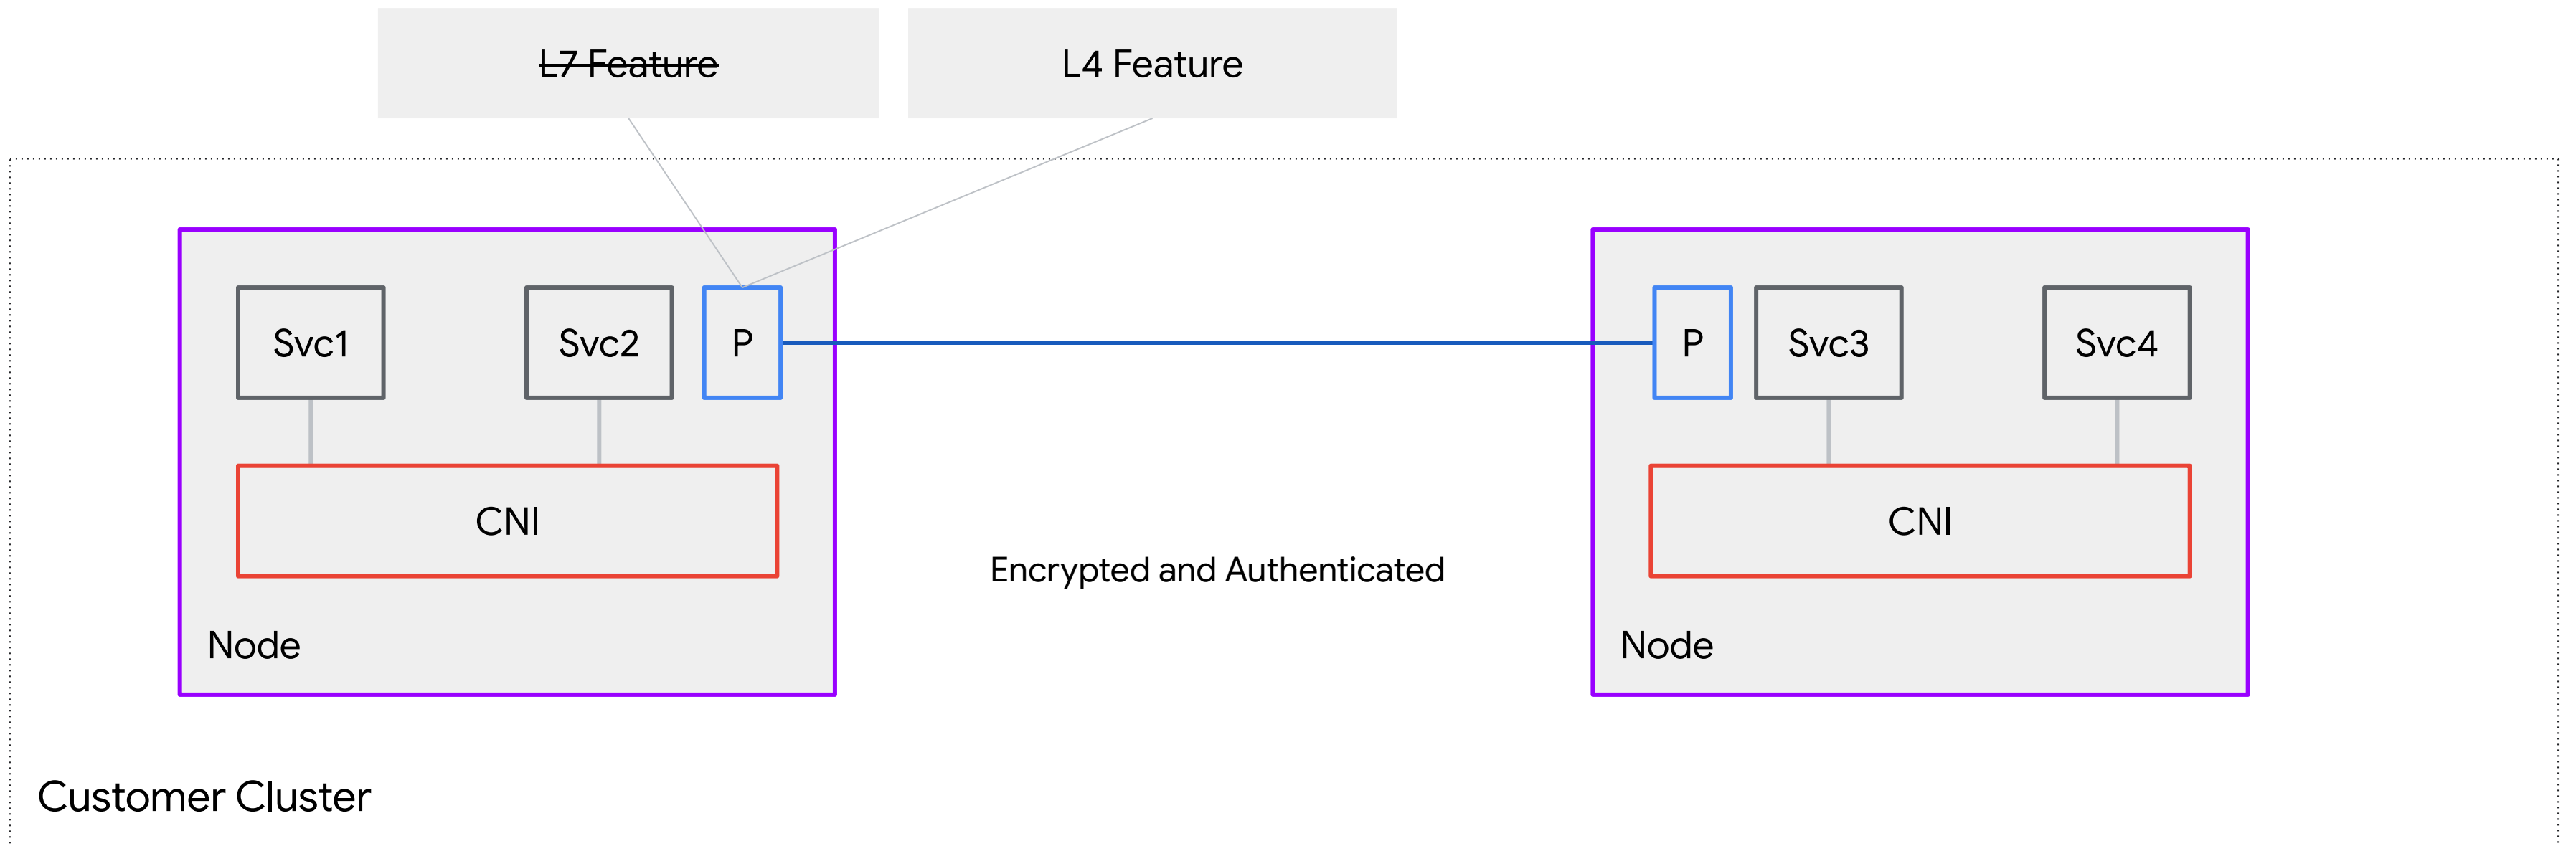


Customer Cluster

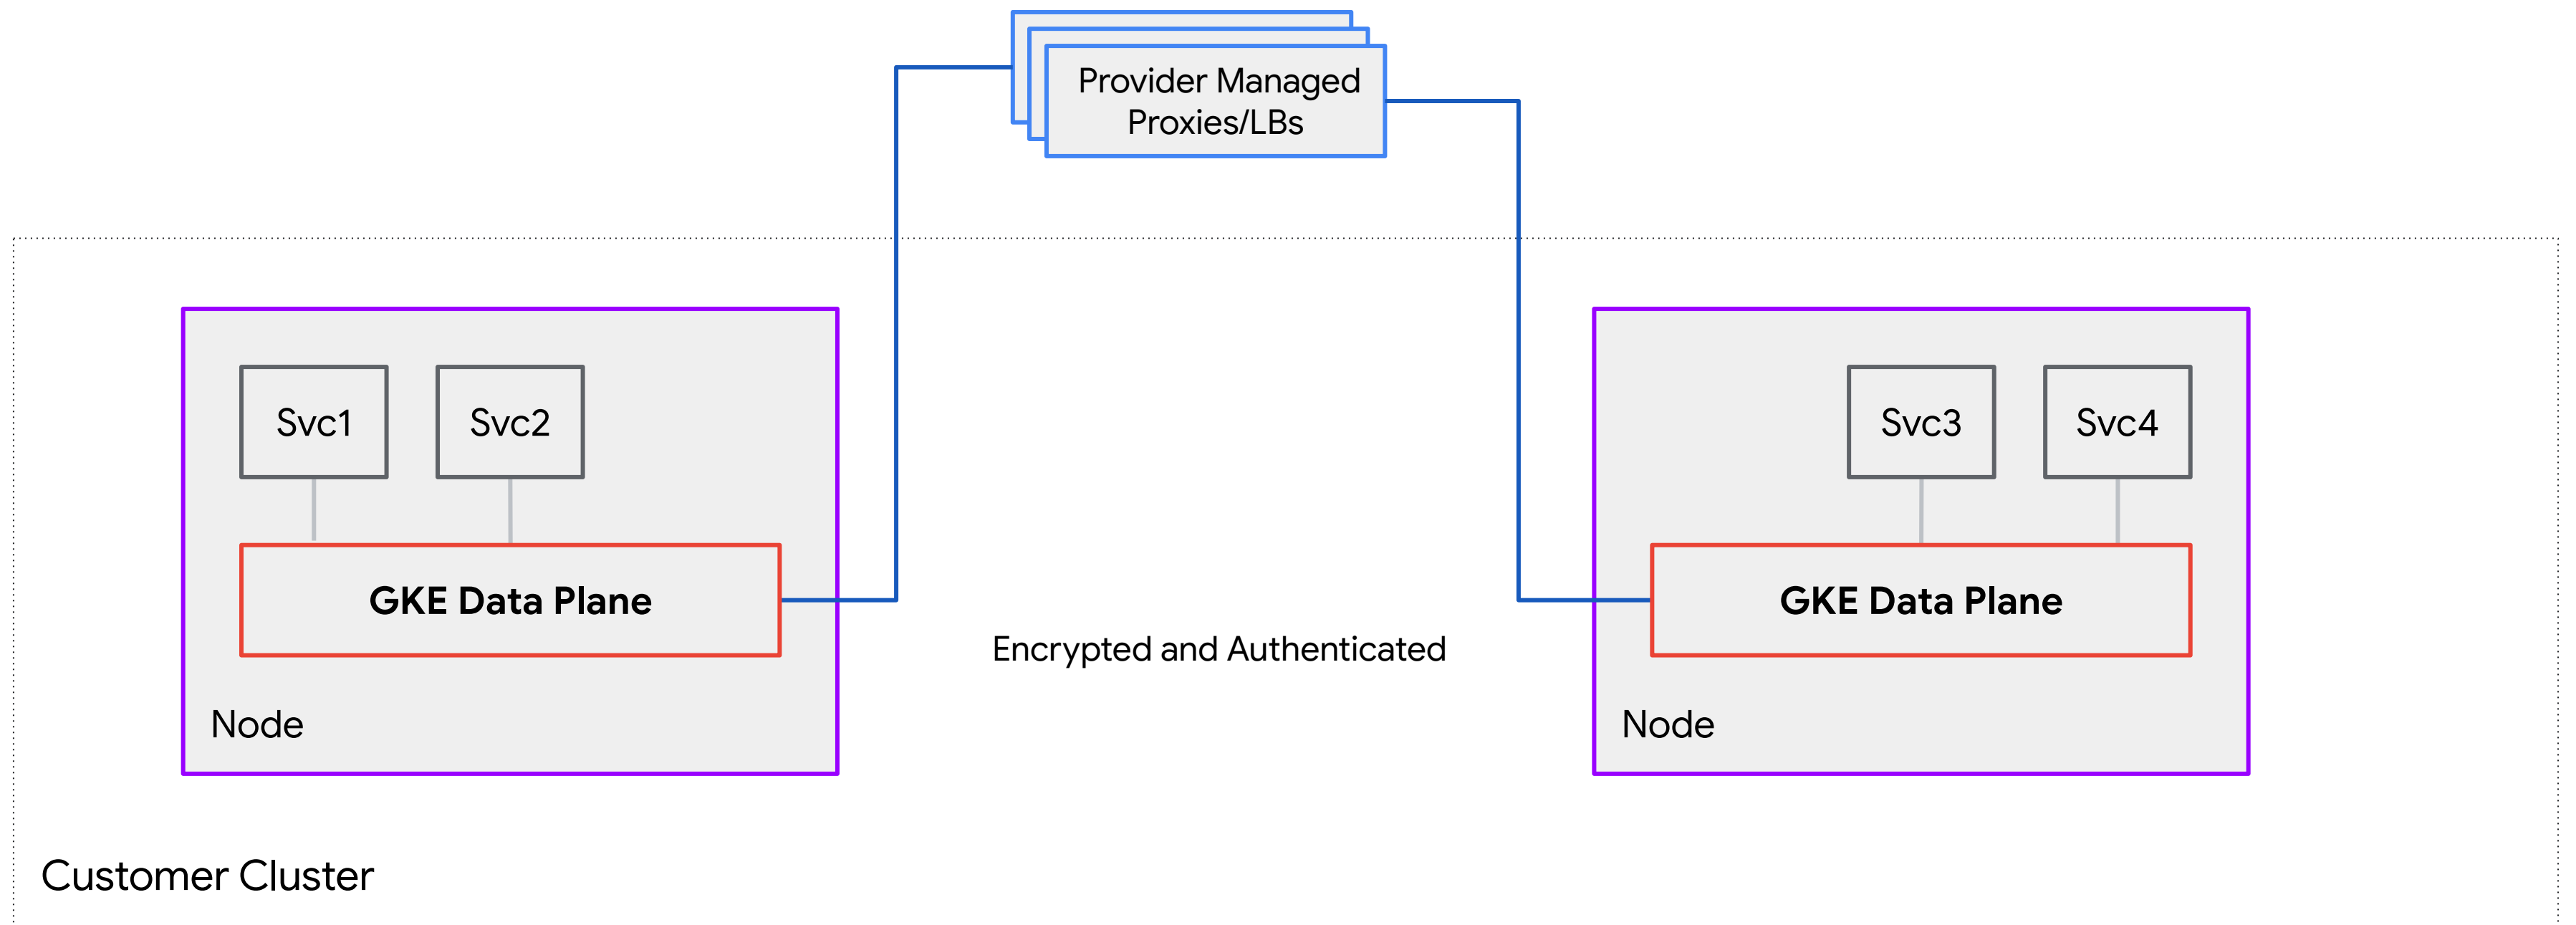
Mesh with sidecars



Layer 4 Ambient Mesh



Ambient Mesh with Provider Managed Mesh



API

API

Istio APIs

Compatible with stable Istio features

API

Istio APIs

Compatible with stable Istio features

Google APIs

GCP native API surface

API

Istio APIs

Compatible with stable Istio features

Google APIs

GCP native API surface

Gateway for Mesh API

K8s vendor neutral API surface
GKE vendor extensions

Data Plane Extensibility

Extensibility with CSM

WASM

Web Assembly - Standard mechanism

WASM As A Service as managed offering

Language Choices

Marketplace integration

Service Callouts

ext_proc callouts from data plane

EnvoyFilters

Break glass use cases: For Envoy only

Prefer first class API, WASM & DEP

Evaluating filters for use on a case-by-case basis

Istio

Istio and CSM

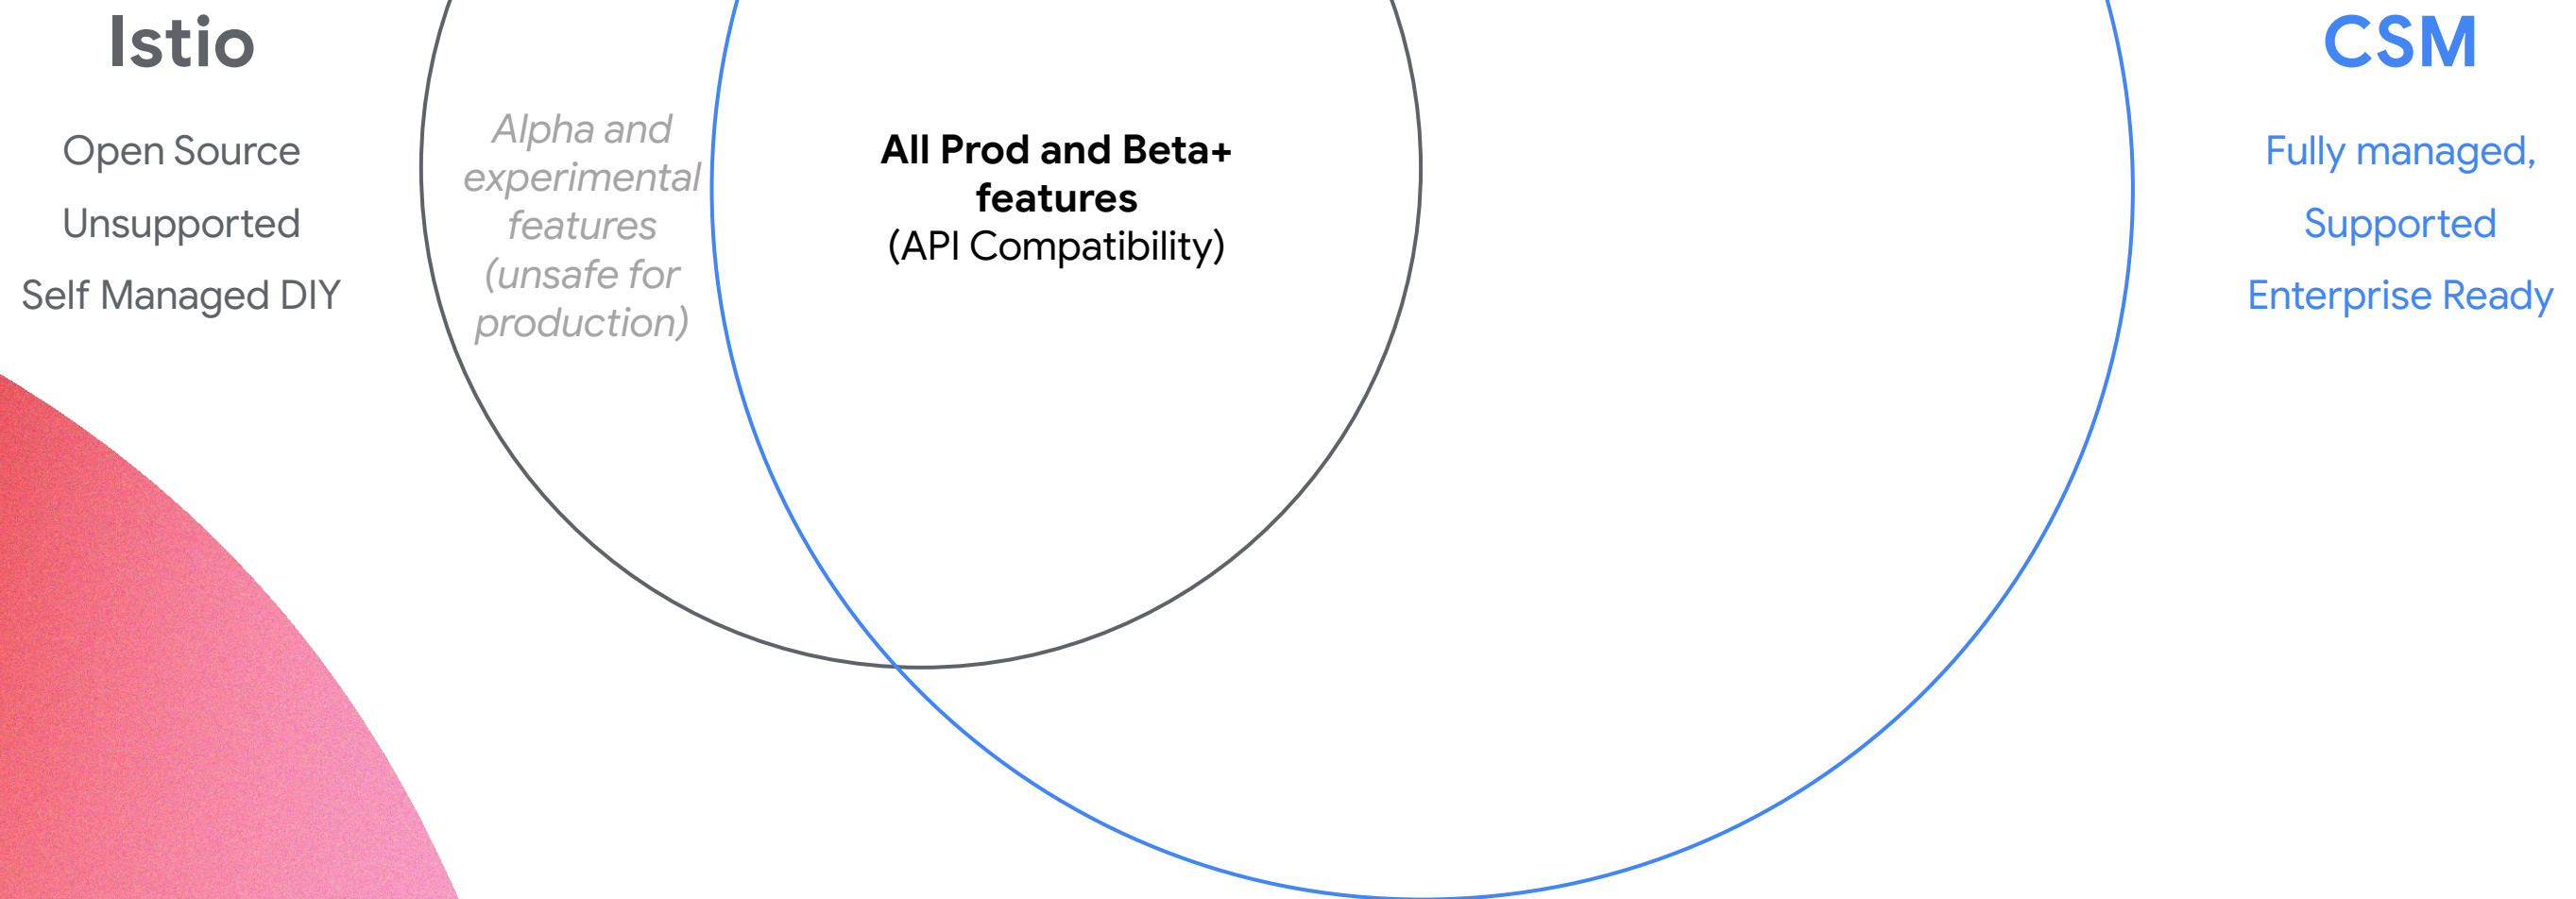
Istio

Open Source
Unsupported
Self Managed DIY

*Alpha and
experimental
features
(unsafe for
production)*

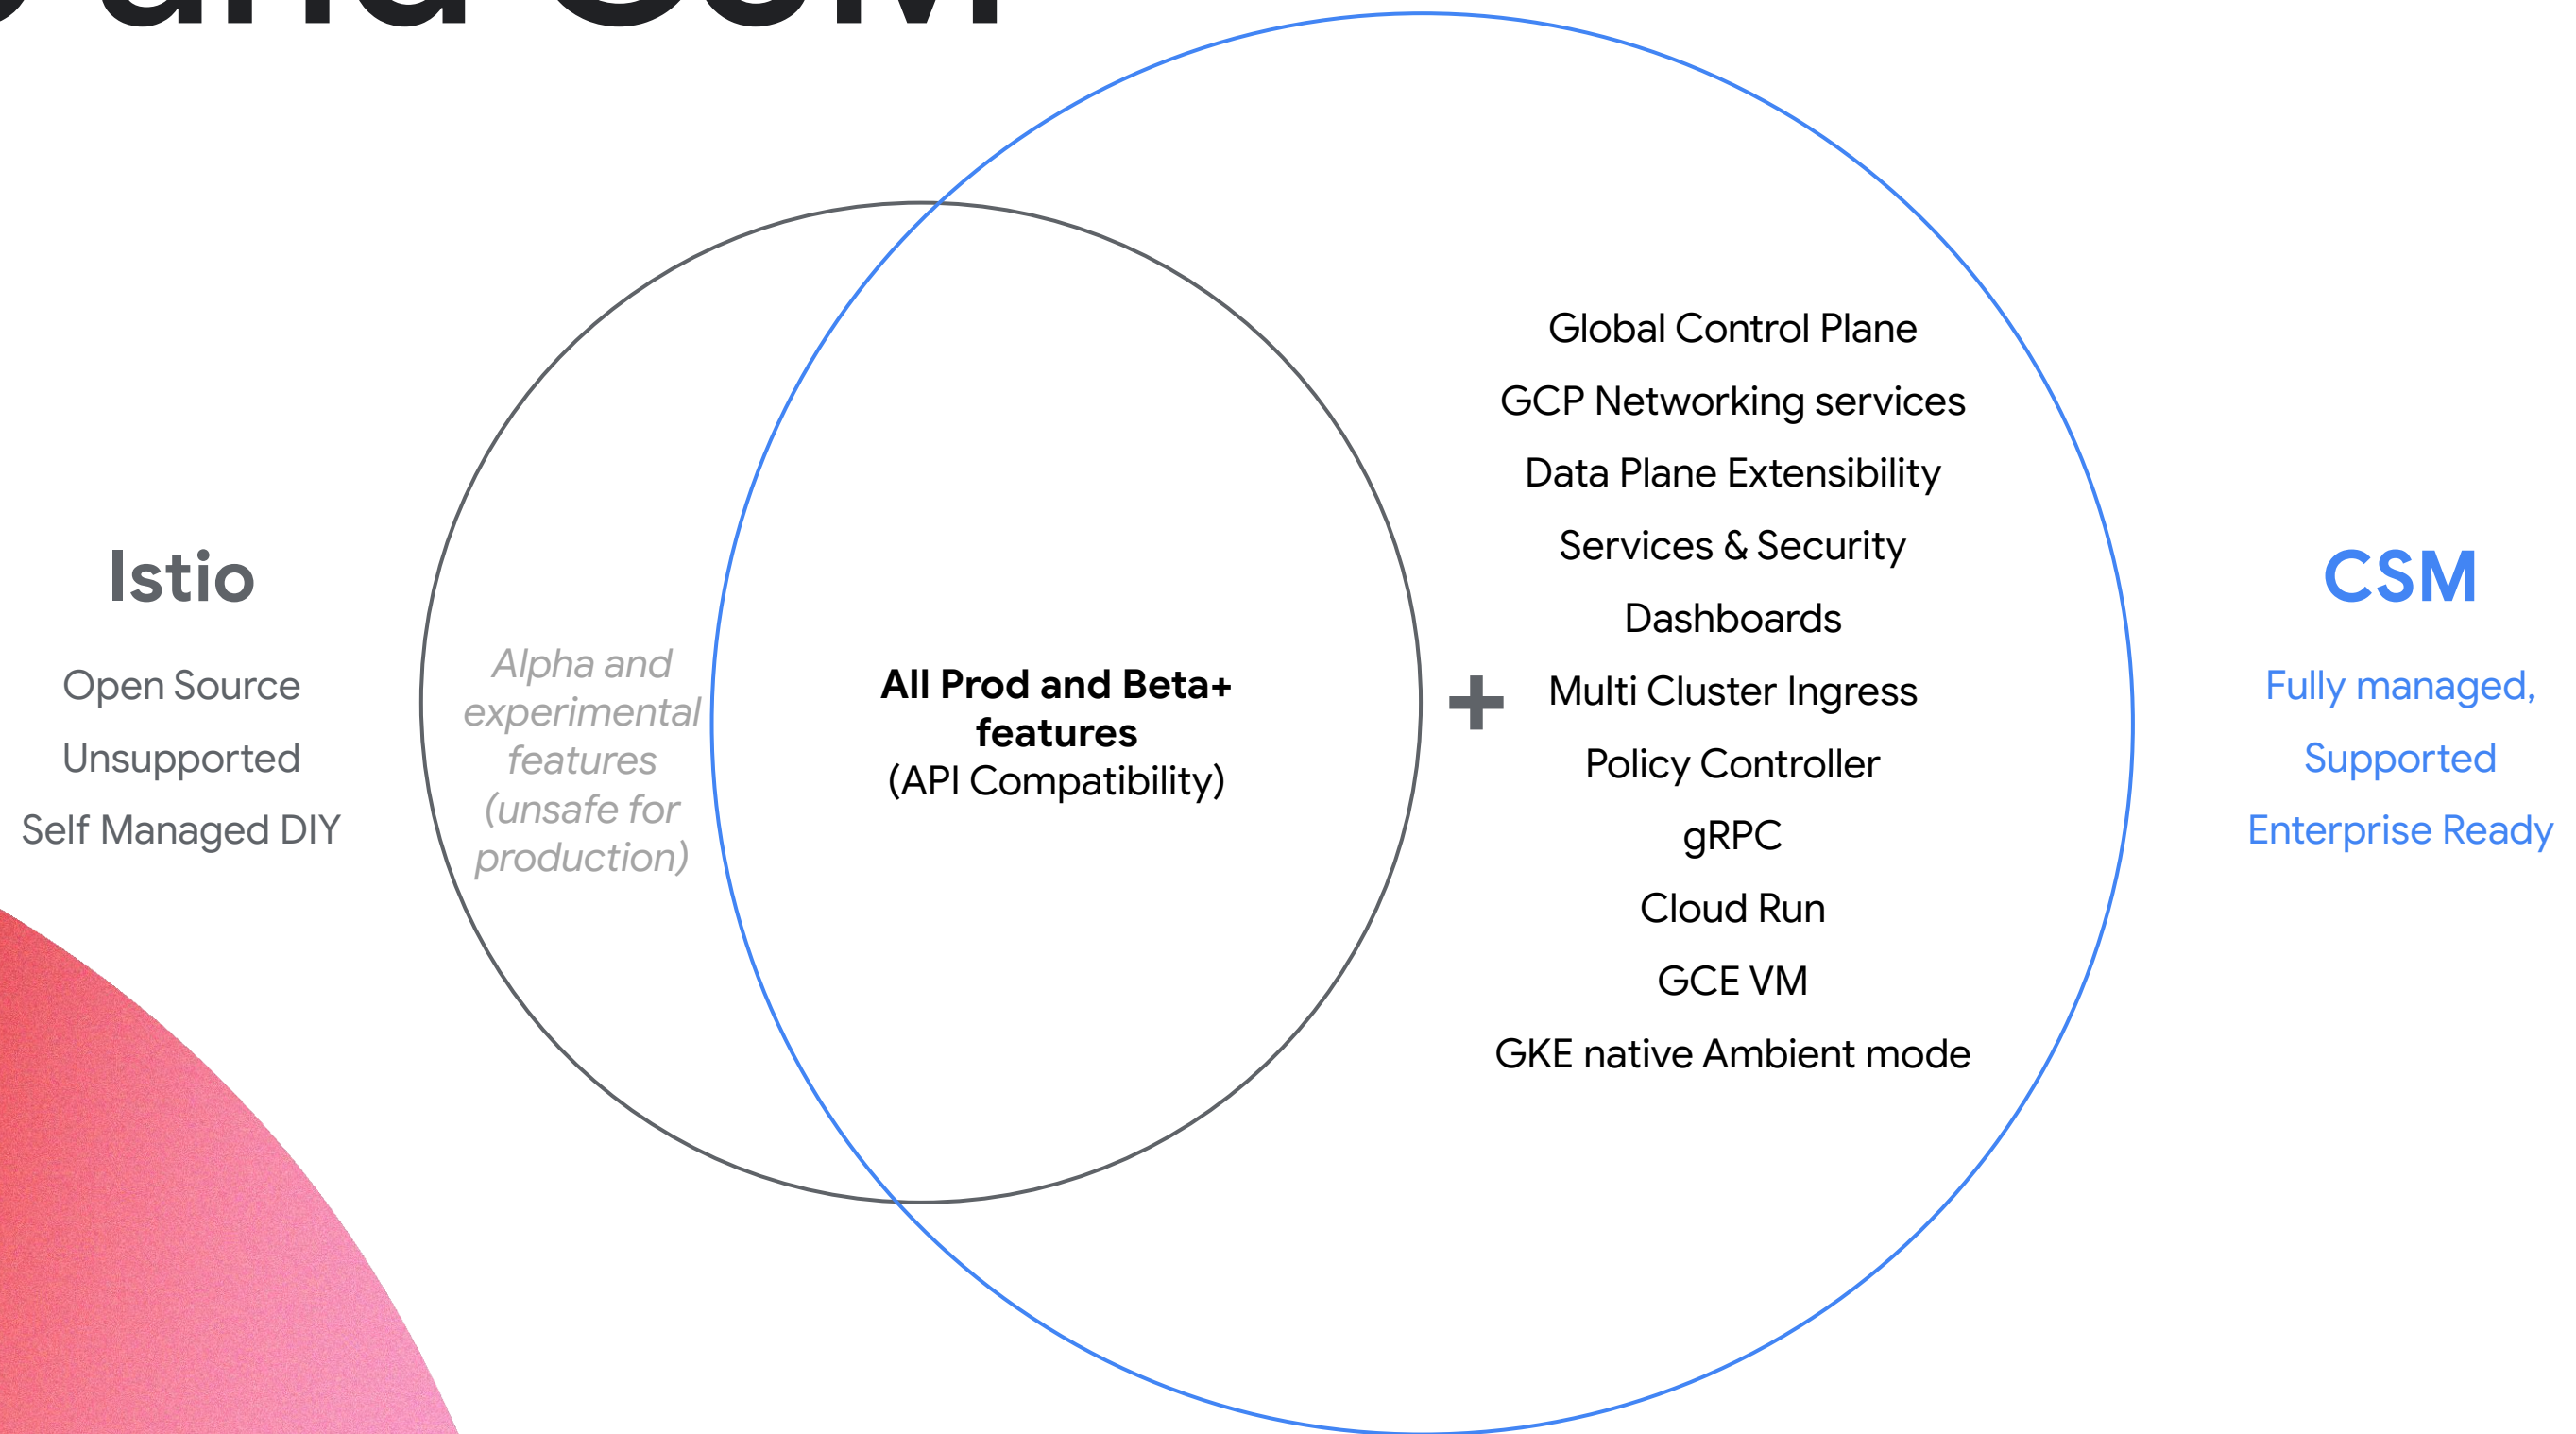
**All Prod and Beta+
features**
(API Compatibility)

Istio and CSM




Proprietary

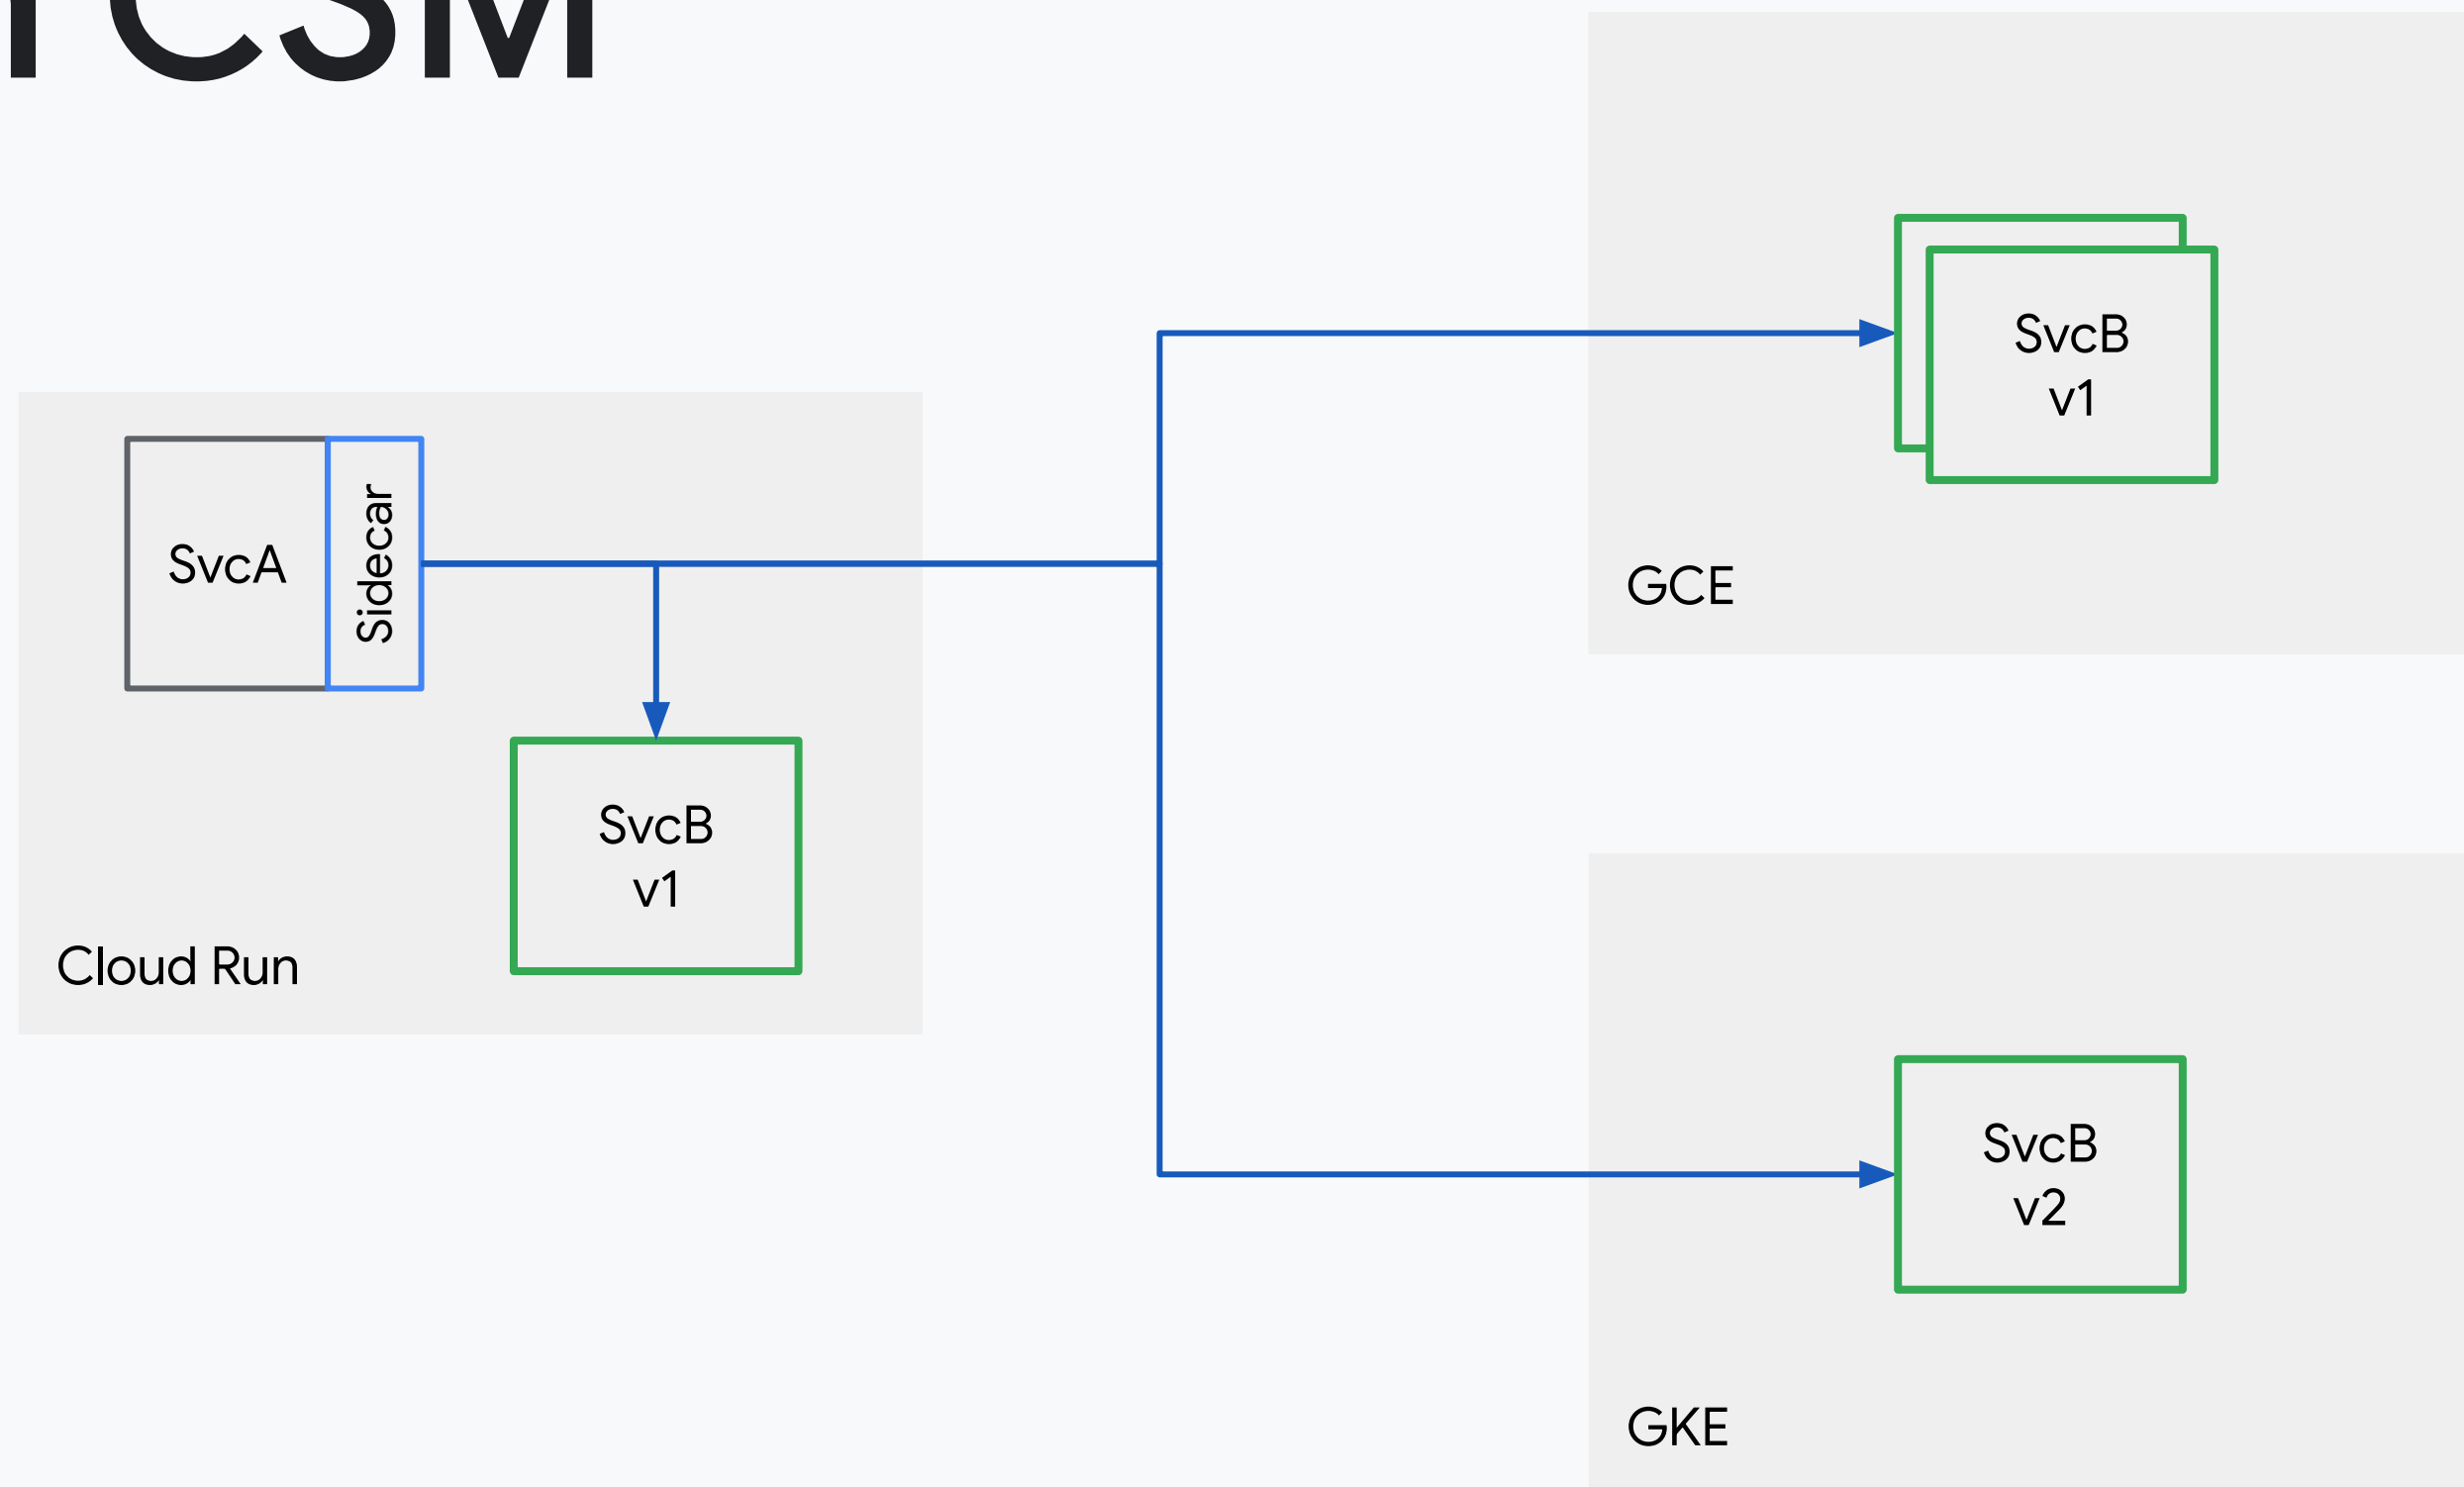
Istio and CSM



CSM Anywhere



Multi runtime architecture with CSM



```
// HTTP route targetting the GCE MIG, Cloud Run Service, and GKE Service.
```

```
resource "google_network_services_http_route" "helloworld" {
  provider = google-beta
  name      = "helloworld-route"
  meshes = [
    google_network_services_mesh.default.id
  ]
  hostnames = [format("hello.%s", var.domain)]
  rules {
    matches {
      full_path_match = "/"
    }
    action {
      url_rewrite {
        path_prefix_rewrite = "/hello"
      }
      # destinations {
      #   service_name = module.mig.backend_service
      # }
      # destinations {
      #   service_name = module.run.backend_service
      # }
      # destinations {
      #   service_name = module.gke.backend_service
      # }
      retry_policy {
        retry_conditions = ["5xx"]
        num_retries = 2
      }
    }
  }
  lifecycle {
    create_before_destroy = true
  }
}
```

```
resource "google_cloud_run_service" "mesh-service" {
  name = "hello-mesh-test"
  location = var.region

  metadata {
    annotations = {
      "run.googleapis.com/launch-stage" = "ALPHA"
    }
  }
}
```

```
-UU-:***- F1 main.tf 42% L66 (Terraform WK counsel ivy FlyC EditorConfig) ---
```

```
. [id=projects/ianmi-test/regions/us-central1/instanceGroupManagers/hello-mig, 10s elapsed]
]
module.mig.google_compute_health_check.http-hello: Destruction complete after 11s
module.mig.module.mig.google_compute_region_instance_group_manager.mig: Still destroying..
. [id=projects/ianmi-test/regions/us-central1/instanceGroupManagers/hello-mig, 20s elapsed]
]
module.mig.module.mig.google_compute_region_instance_group_manager.mig: Still destroying..
. [id=projects/ianmi-test/regions/us-central1/instanceGroupManagers/hello-mig, 30s elapsed]
]
module.mig.module.mig.google_compute_region_instance_group_manager.mig: Still destroying..
. [id=projects/ianmi-test/regions/us-central1/instanceGroupManagers/hello-mig, 40s elapsed]
]
module.mig.module.mig.google_compute_region_instance_group_manager.mig: Destruction complete after 41s
module.mig.module.mig_template.google_compute_instance_template.tpl: Destroying... [id=projects/ianmi-test/global/instanceTemplates/hello-mig-20240310190842353400000001]
module.mig.module.mig_template.google_compute_instance_template.tpl: Still destroying... [id=projects/ianmi-test/global/instanceTemplates/hello-mig-20240310190842353400000001, 10s elapsed]
module.mig.module.mig_template.google_compute_instance_template.tpl: Destruction complete after 11s
```

```
Apply complete! Resources: 0 added, 1 changed, 12 destroyed.
```

```
[hi on] ianmi@ianmi:~/git/cloud-run-mesh-samples/http-routing$
```

```
-n, --interval <secs> seconds to wait between updates
-p, --precise          attempt run command in precise intervals
-r, --no-rerun        do not rerun program on window resize
-t, --no-title        turn off header
-w, --no-wrap         turn off line wrapping
-x, --exec            pass command to exec instead of "sh -c"
```

```
-h, --help            display this help and exit
-v, --version        output version information and exit
```

```
For more details see watch(1).
```

```
[hi on] ianmi@ianmi:~$ watch -n 1 curl -v http://127.0.0.1:8080/hello
```

```
[hi on] ianmi@ianmi:~$ man lynx
```

```
[hi on] ianmi@ianmi:~$ lynx -dump http://127.0.0.1:8080/hello
```

```
Hello version: v1, instance: default-xqx9
```

```
[hi on] ianmi@ianmi:~$ lynx -dump -error_file=/dev/stderr http://127.0.0.1:8080/hello
```

```
URL=http://127.0.0.1:8080/hello (GET)
```

```
STATUS=HTTP/1.0 200 OK
```

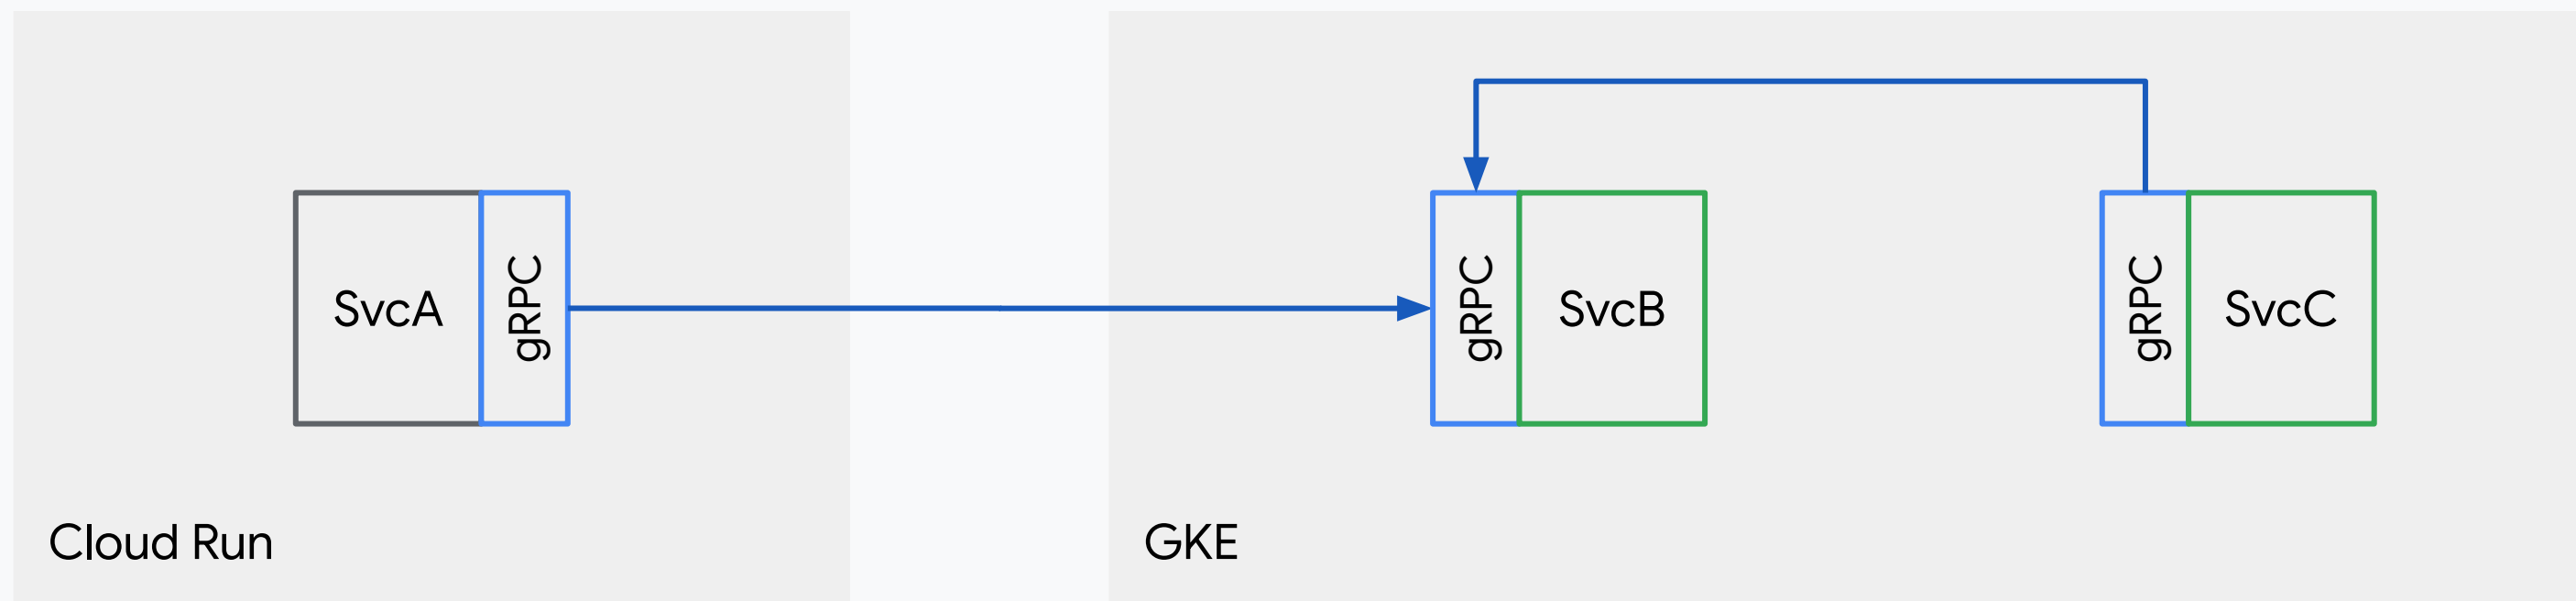
```
Hello version: v1, instance: default-xtp8
```

```
[hi on] ianmi@ianmi:~$ watch -n 1 lynx -dump -error_file=/dev/stderr http://127.0.0.1:8080
```

```
/hello
```

```
[hi on] ianmi@ianmi:~$
```

gRPC data plane with Gateway API



```
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: echo
  namespace: default
spec:
  selector:
    matchLabels:
      app: echo
  replicas: 1
  template:
    metadata:
      labels:
        app: echo
    spec:
      containers:
        - name: echo
          image: gcr.io/istio-testing/app:latest
          args:
            - --port=8080
            - --grpc=50051
          ports:
            - containerPort: 50051
              name: grpc
          livenessProbe:
            initialDelaySeconds: 1
            grpc:
              port: 50051
---
apiVersion: v1
kind: Service
metadata:
  name: echo-grpc
  namespace: default
spec:
  ports:
    - name: grpc
      port: 50051
      targetPort: 50051
  selector:
    app: echo
---
apiVersion: gateway.networking.k8s.io/v1alpha2
kind: GRPCRoute
metadata:
  name: echo
spec:
  parentRefs:
    - name: echo-grpc
      kind: Service
      group: ""
  rules:
    - backendRefs:
        - name: echo-grpc
          port: 50051
---
-UU-:--- F1 grpc deployment.yaml<gateway-api> Top L40 Git:main (YAML ARev WK counsel ivy FlyC EditorConfig) -----
Wrote /usr/local/google/home/ianmi/git/cloud-run-mesh-samples/gateway-api/grpc_deployment.yaml
[mesh-demo0:~]$
```

```
[hi on] ianmi@ianmi:~$
```

CSM Big Picture


More Runtimes and Infrastructure

GKE, GCE, Cloud Run

Single, Shared, Peered VPC or
multi-network support

Istiod (Federated by TD) for off
GCP infrastructure
(multi-cloud and hybrid cloud)

Multi data plane support
(Envoy, gRPC, Cilium, xLBs)



CSM Big Picture

More Runtimes and Infrastructure

GKE, GCE, Cloud Run

Single, Shared, Peered VPC or
multi-network support

Istiod (Federated by TD) for off
GCP infrastructure
(multi-cloud and hybrid cloud)

Multi data plane support
(Envoy, gRPC, Cilium, xLBs)

Better GCP Ecosystem integration

App Hub


Data Plane Extensibility
Platform (DEP)

Managed LBs

HCaaS (Health Checking)

PSC

Cross cloud networking



CSM Big Picture

More Runtimes and Infrastructure

GKE, GCE, Cloud Run

Single, Shared, Peered VPC or
multi-network support

Istiod (Federated by TD) for off
GCP infrastructure
(multi-cloud and hybrid cloud)

Multi data plane support
(Envoy, gRPC, Cilium, xLBs)

Better GCP Ecosystem integration

App Hub

Data Plane Extensibility
Platform (DEP)

Managed LBs

HCaaS (Health Checking)

PSC

Cross cloud networking

Focus on Services not mesh

Focus on consuming features, not
managing a mesh product

AppHub integration

Intelligent mesh (always on)

Ready to build what's next?

Tap into **special offers** designed to help you **implement what you learned** at Google Cloud Next.

Scan the code to receive personalized guidance from one of our experts.




Or visit g.co/next/24offers

Continue your learning journey!




**DEV 205 -
Cloud Run: What's New**

April 10 @ 3:30



**SEC305 -
Certifiably secure:
Identities for workloads
and devices with
Certificate Authority**

April 11 @ 1:30 PM



**The future of platform
engineering is
application-specific**

April 09 @ 11 AM

Thank you

