



KubeCon



CloudNativeCon

Europe 2018

# Enabling Multi-Cloud with Istio

Stretching an Istio service mesh between Public & Private Clouds

John Joyce Cisco Principal Engineer

Robert Li Cisco Software Engineer



# Agenda



KubeCon



CloudNativeCon

Europe 2018

- Introduction
- Multicluster Architecture
- Deployment Details
- Deployment in Action
- Mesh Routing Demo
- Future Improvements
- References
- Q&A

# Introduction

# Extending an Application across Multiple Clouds

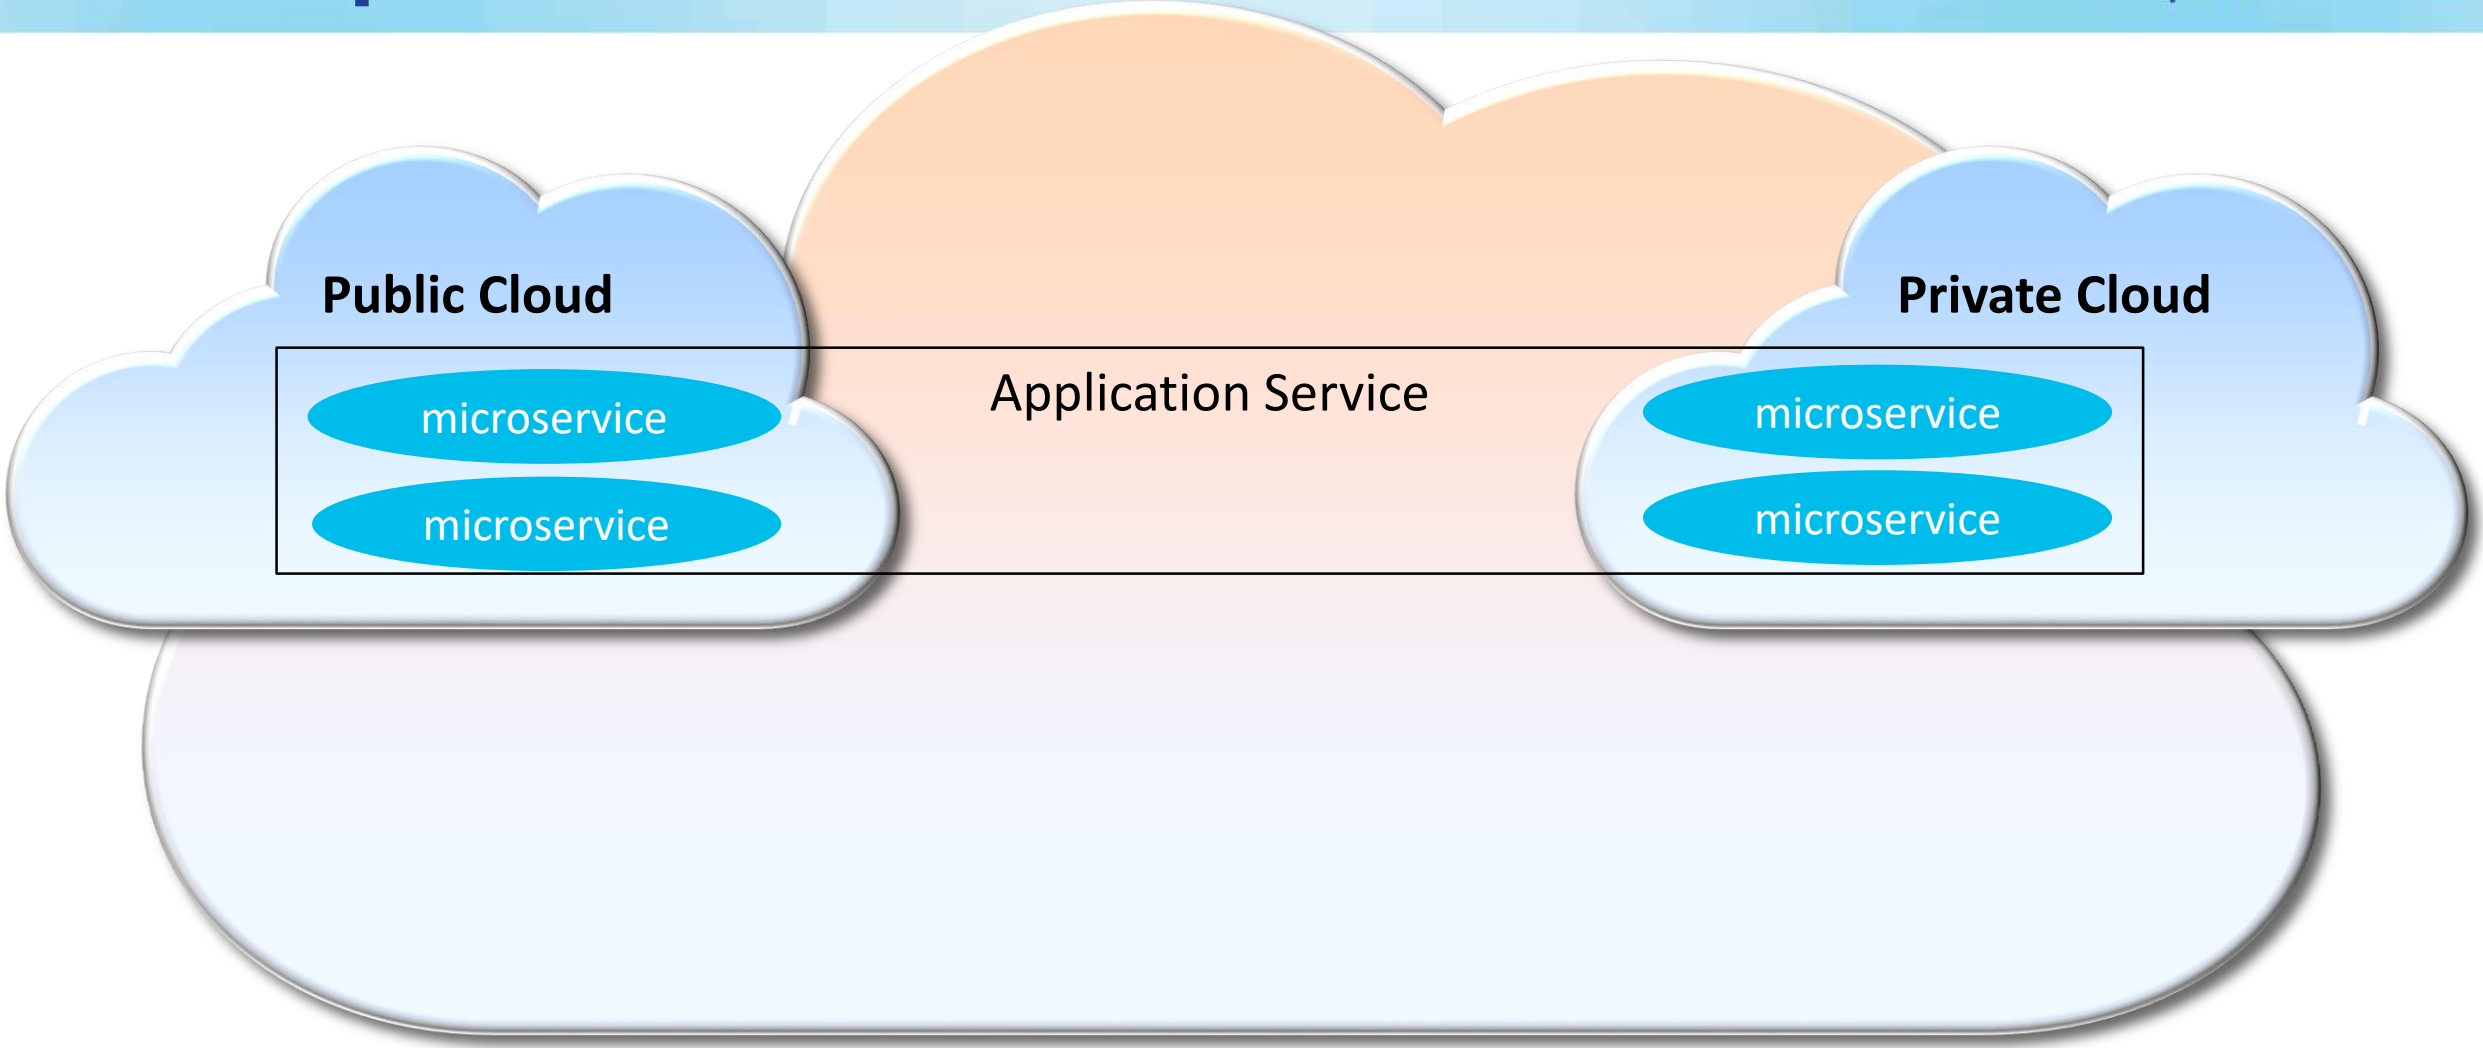


KubeCon



CloudNativeCon

Europe 2018



# Stretching Istio's Service Mesh Across Multiple Clouds

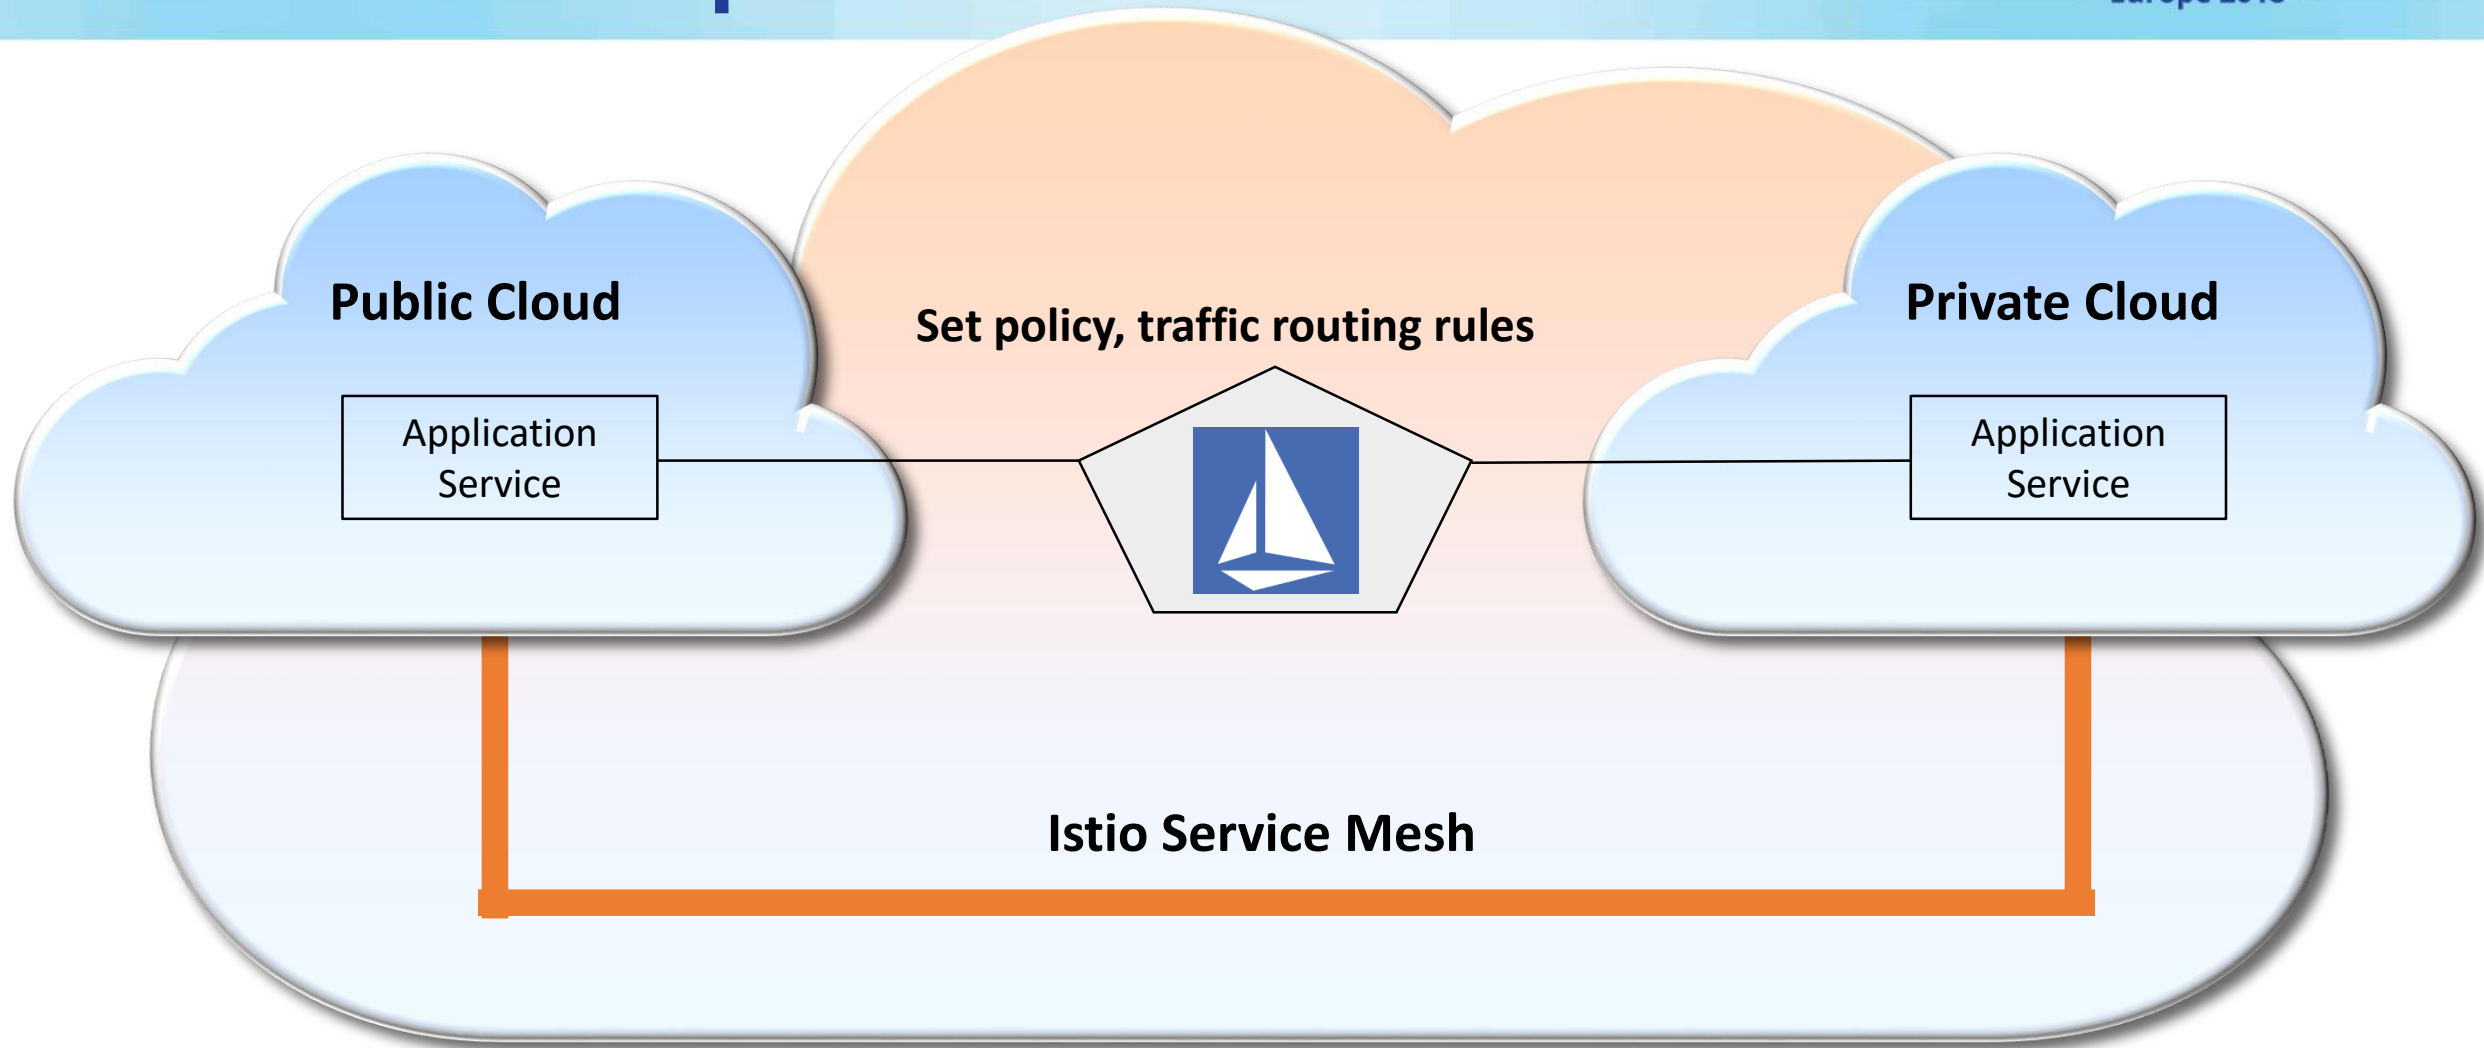


KubeCon

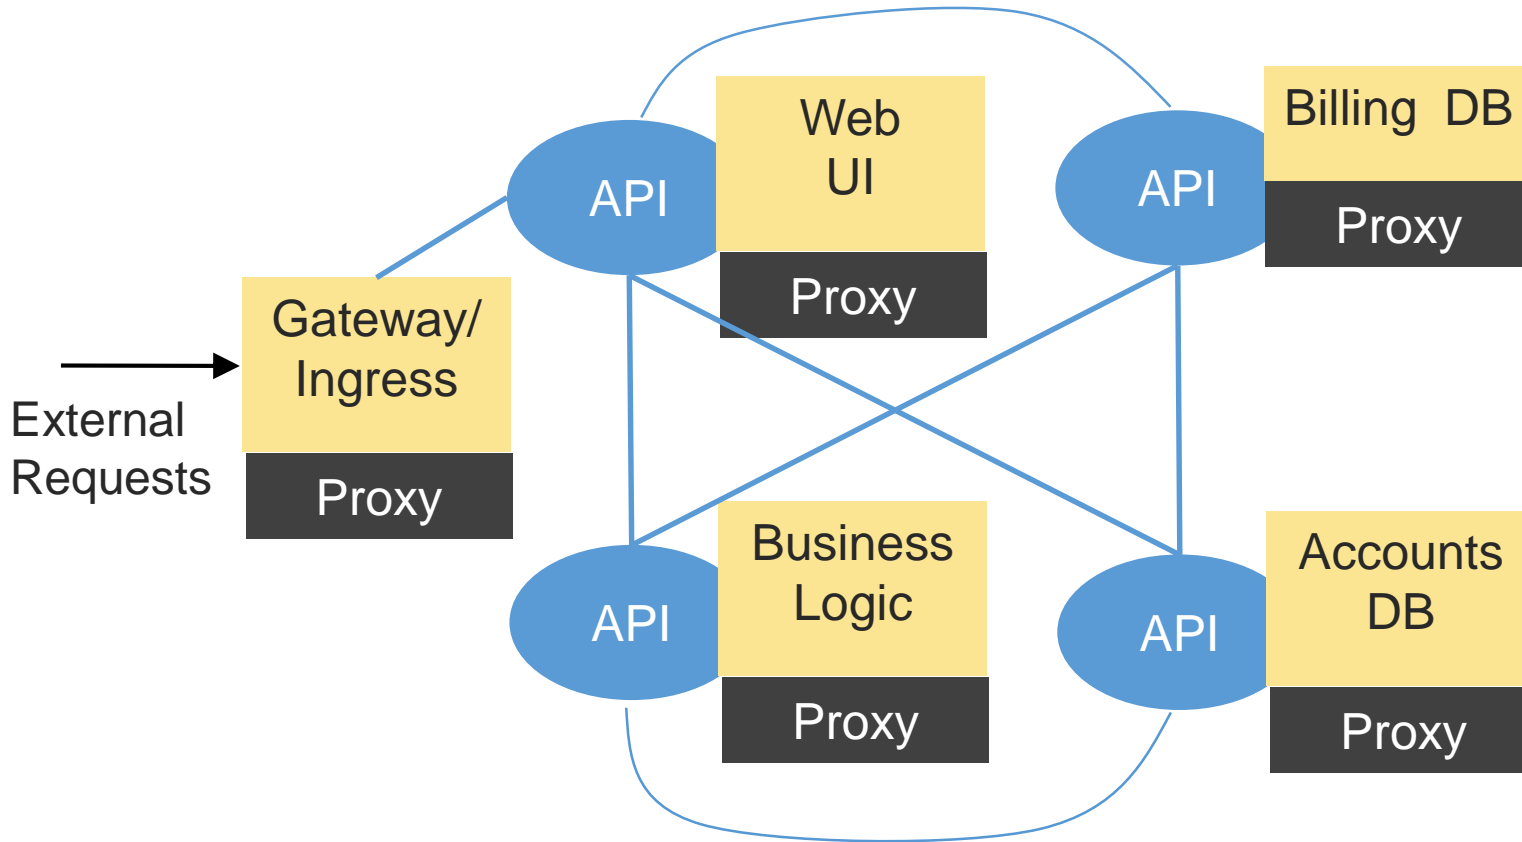


CloudNativeCon

Europe 2018



# What is a Service Mesh



- Infrastructure layer for service-to-service communication
- A mesh of proxies
- Proxies injected as sidecars
- Supports numerous protocols (HTTP 1/2, gRPC, TCP, UDP)
- Can inspect API transactions at Layer 7 or layer 3/4.
- Intelligent routing rules can be applied between endpoints

Istio's Multicluster enhancement allows cloud boundaries anywhere in above mesh

# Istio Architecture

Source: <https://istio.io/docs/concepts/what-is-istio/overview.html>

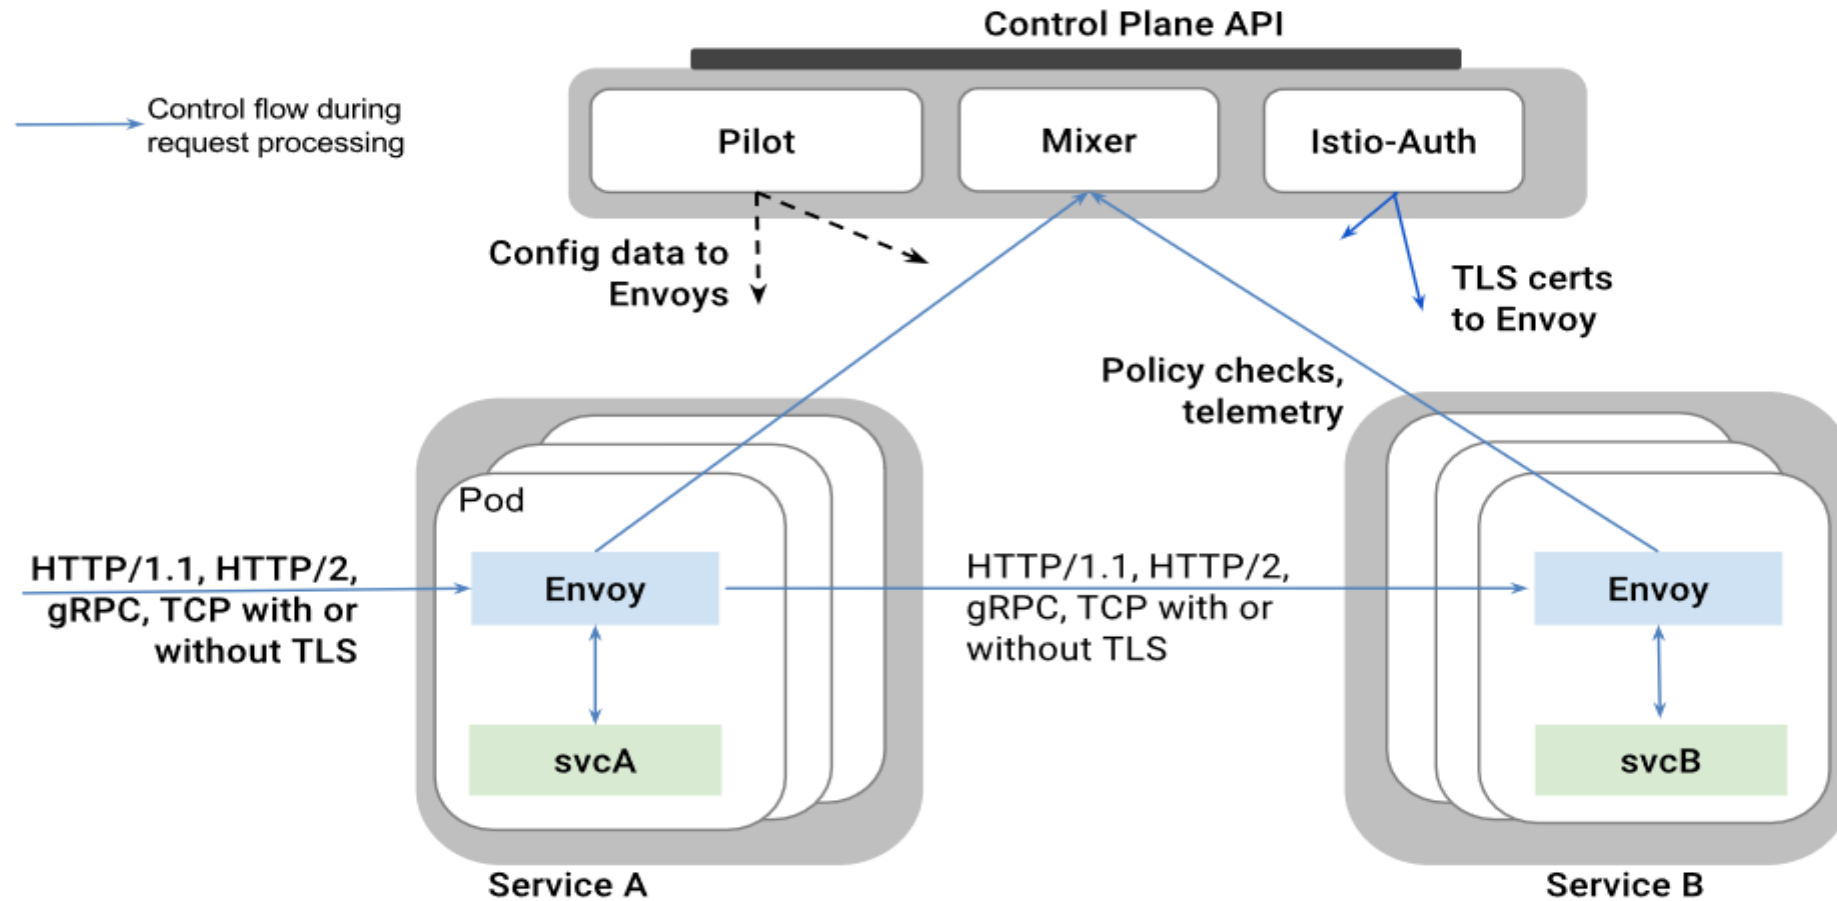


KubeCon



CloudNativeCon

Europe 2018



# Multicluster Architecture

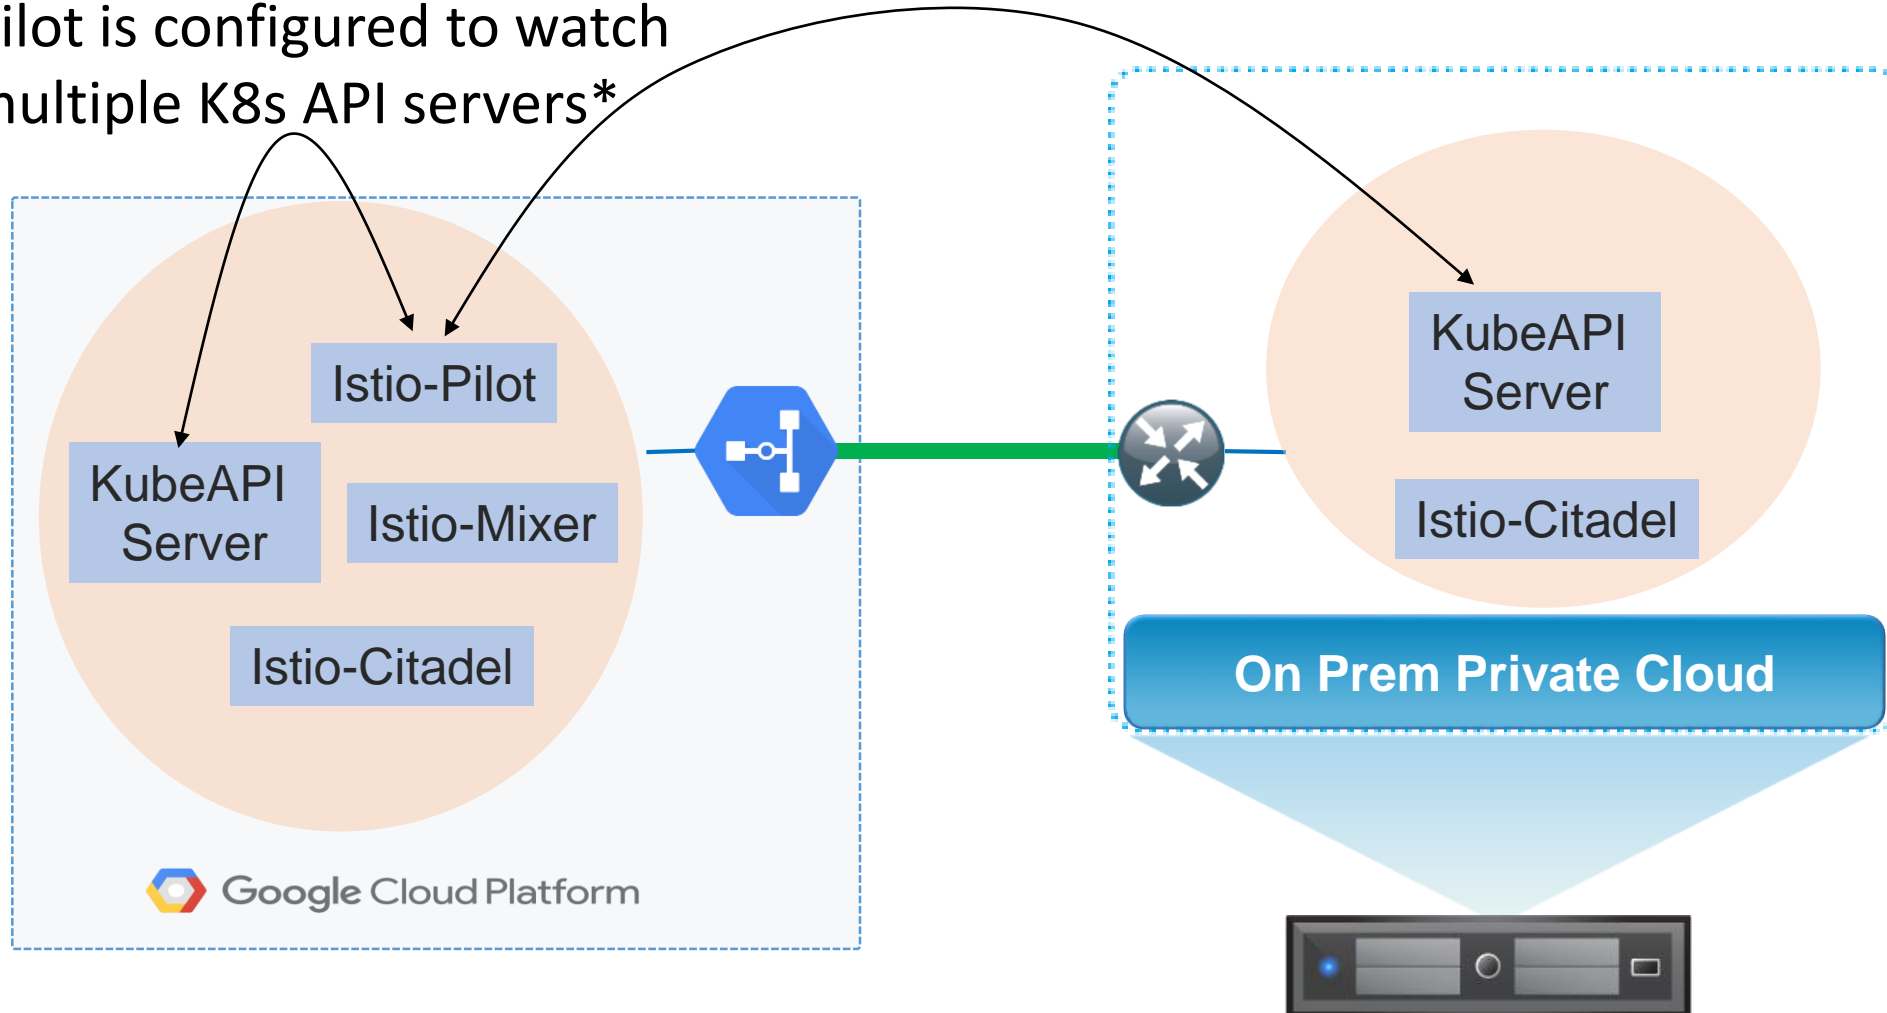


# Service and Endpoint Discovery



Europe 2018

Pilot is configured to watch multiple K8s API servers\*



K8s Clusters

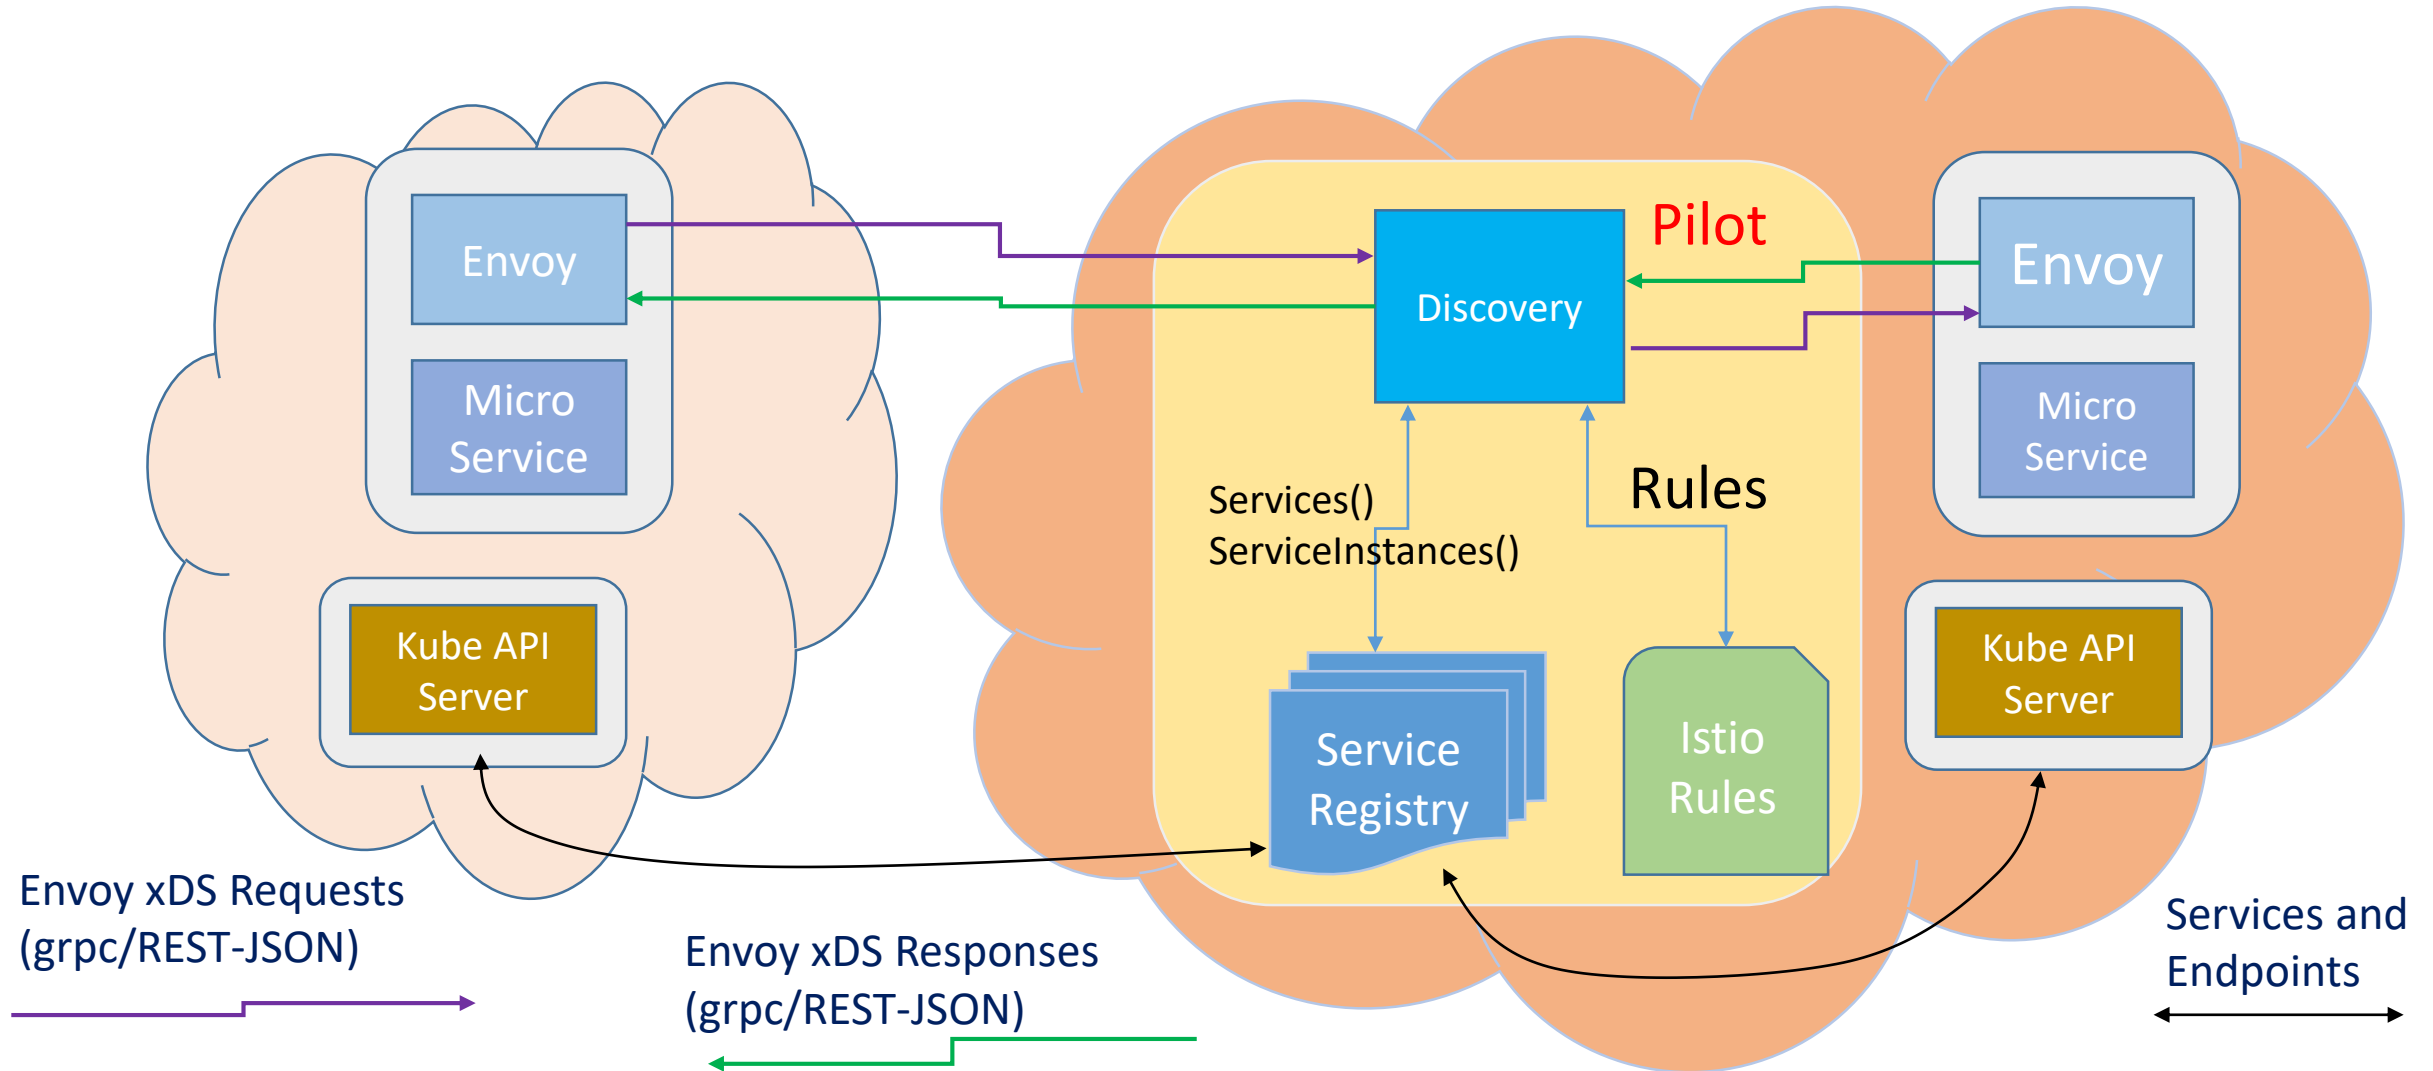
↔  
Controller created in Pilot for each cluster which gathers all service & endpoint data

\* Configuration via file based clusterregistry format

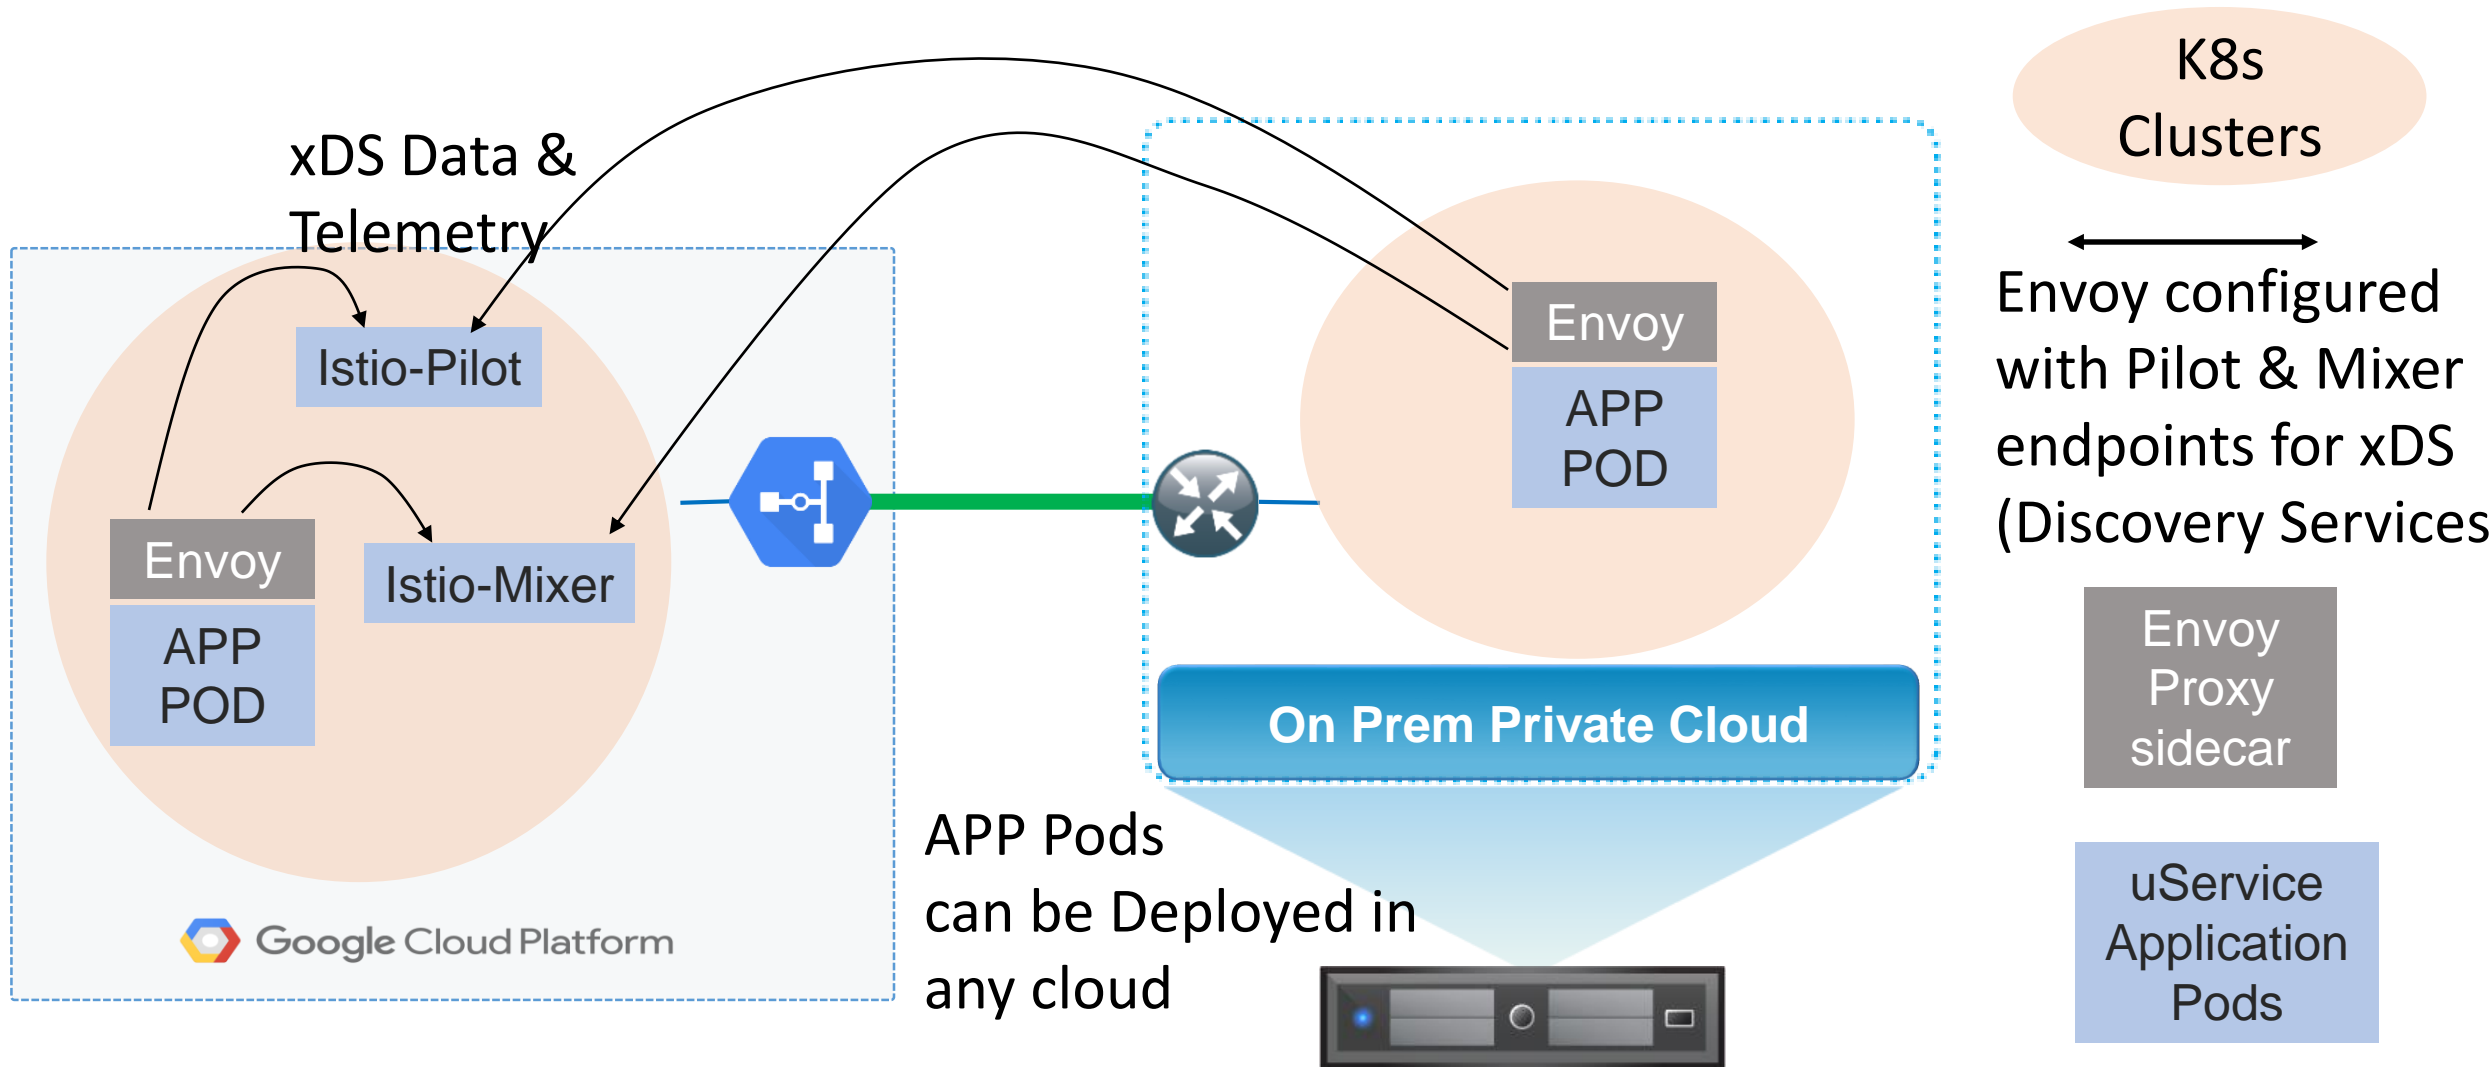
# Envoy xDS propagation



Europe 2018



# Creating the Stretched Mesh



# istio-citadel with same certificates



KubeCon



CloudNativeCon

Europe 2018

On-prem

Istio-  
citadel

- Generate the same ca-cert.pem, ca-key.pem, cert-chain.pem, and root-cert.pem in both clusters
- Create a k8s secret called cacerts from those files
- Add a secret volume from cacerts
- Mount the volume to /etc/cacerts
- Start istio-citadel with the following args:
  - ❖ `--self-signed-ca=false`
  - ❖ `--signing-cert=/etc/cacerts/ca-cert.pem`
  - ❖ `--signing-key=/etc/cacerts/ca-key.pem`
  - ❖ `--root-cert=/etc/cacerts/root-cert.pem`
  - ❖ `--cert-chain=/etc/cacerts/cert-chain.pem`

GKE

Istio-  
citadel

# Deployment Details

# Environment Requirements



KubeCon



CloudNativeCon

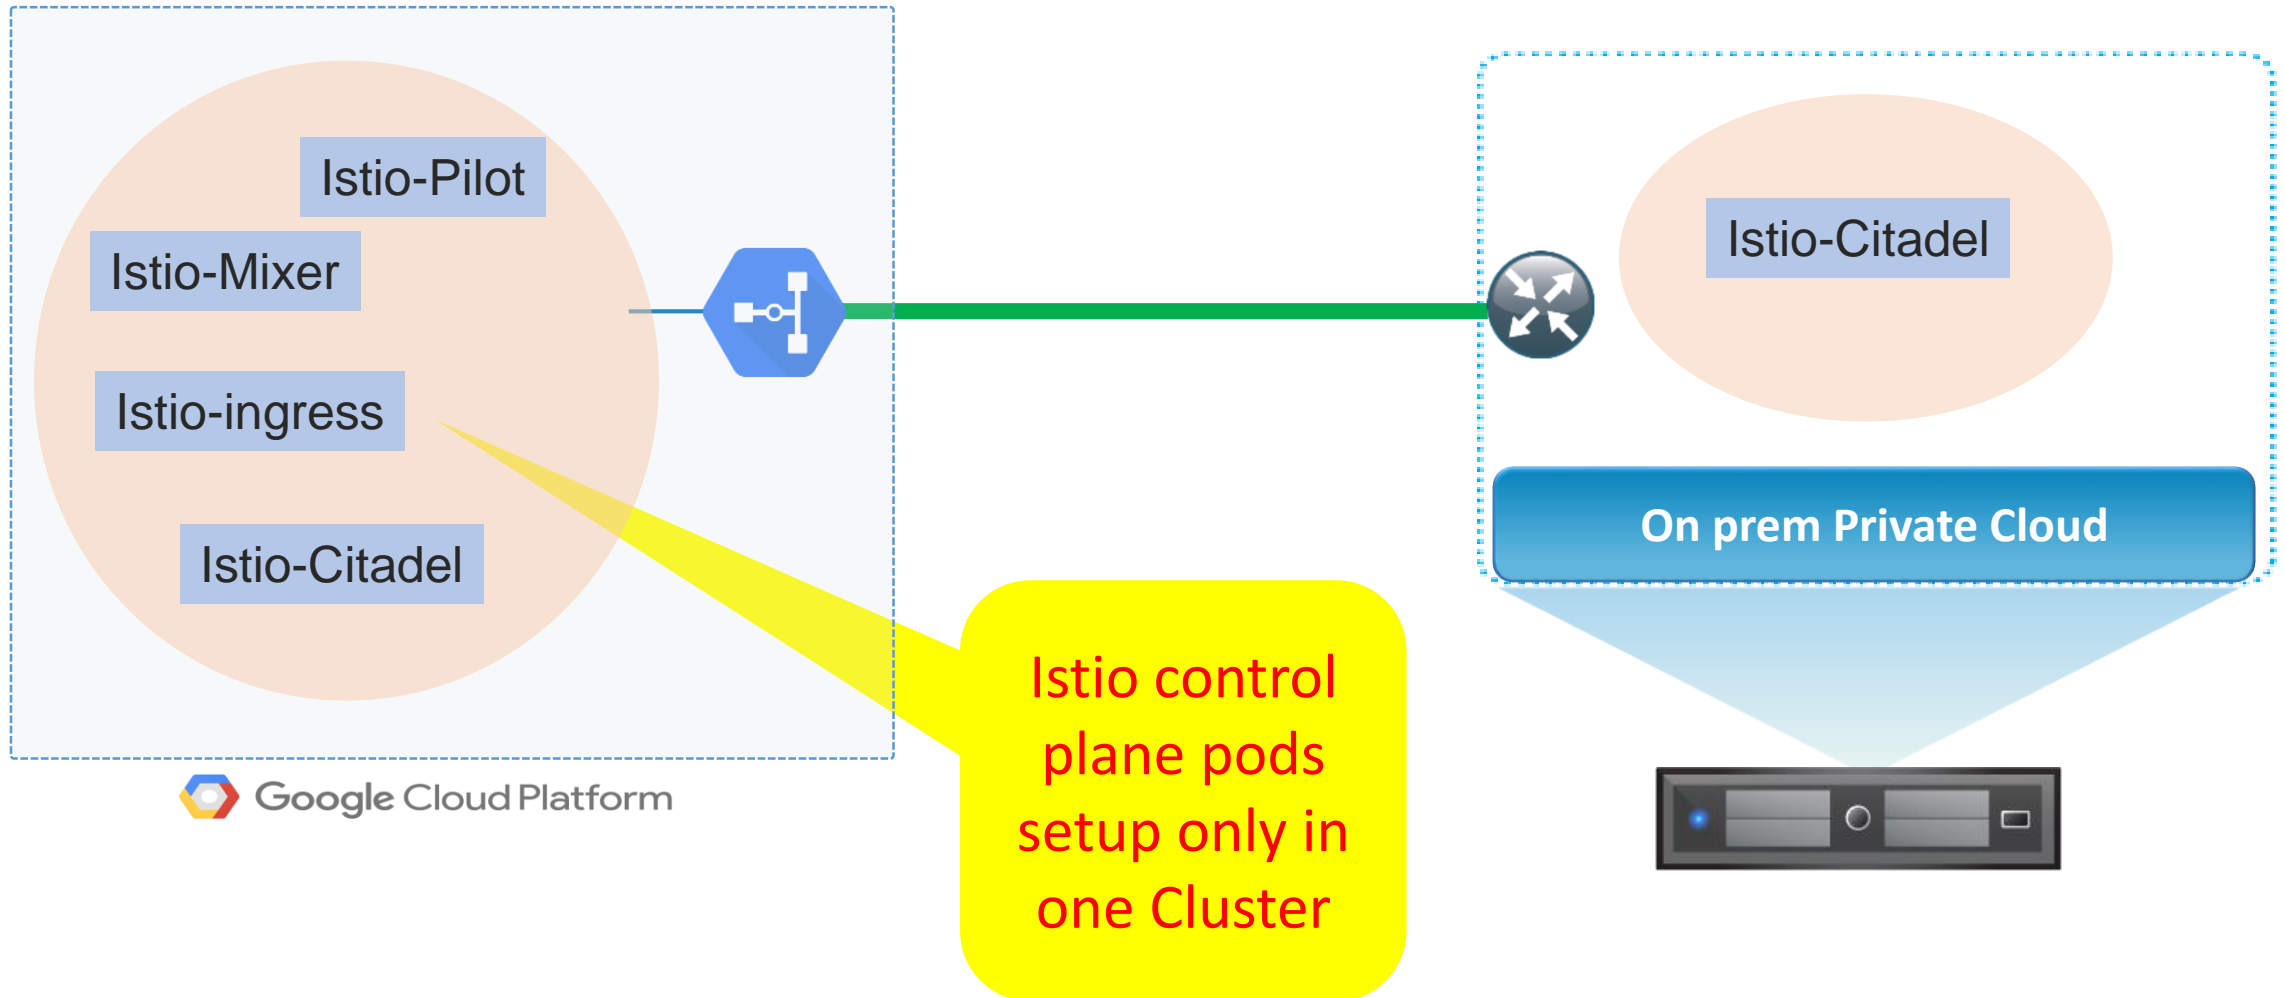
Europe 2018

- Must have IP reachability between the clusters
  - For Istio control plane interactions between Pilot, Mixer and Citadel and Envoy Sidecars
  - For application pods to reach each other
- Solution doesn't dictate a certain approach to achieve this reachability but generally a VPN would be needed based on current capabilities.
- Application relies on Kube-DNS to resolve service naming so special attention is needed during app deployment:
  - Use of headless or selectorless services typically required
  - Could be designated to a special purpose DNS server

# Istio Control Plane Deployment



Europe 2018



# Bookinfo Application Microservice Architecture

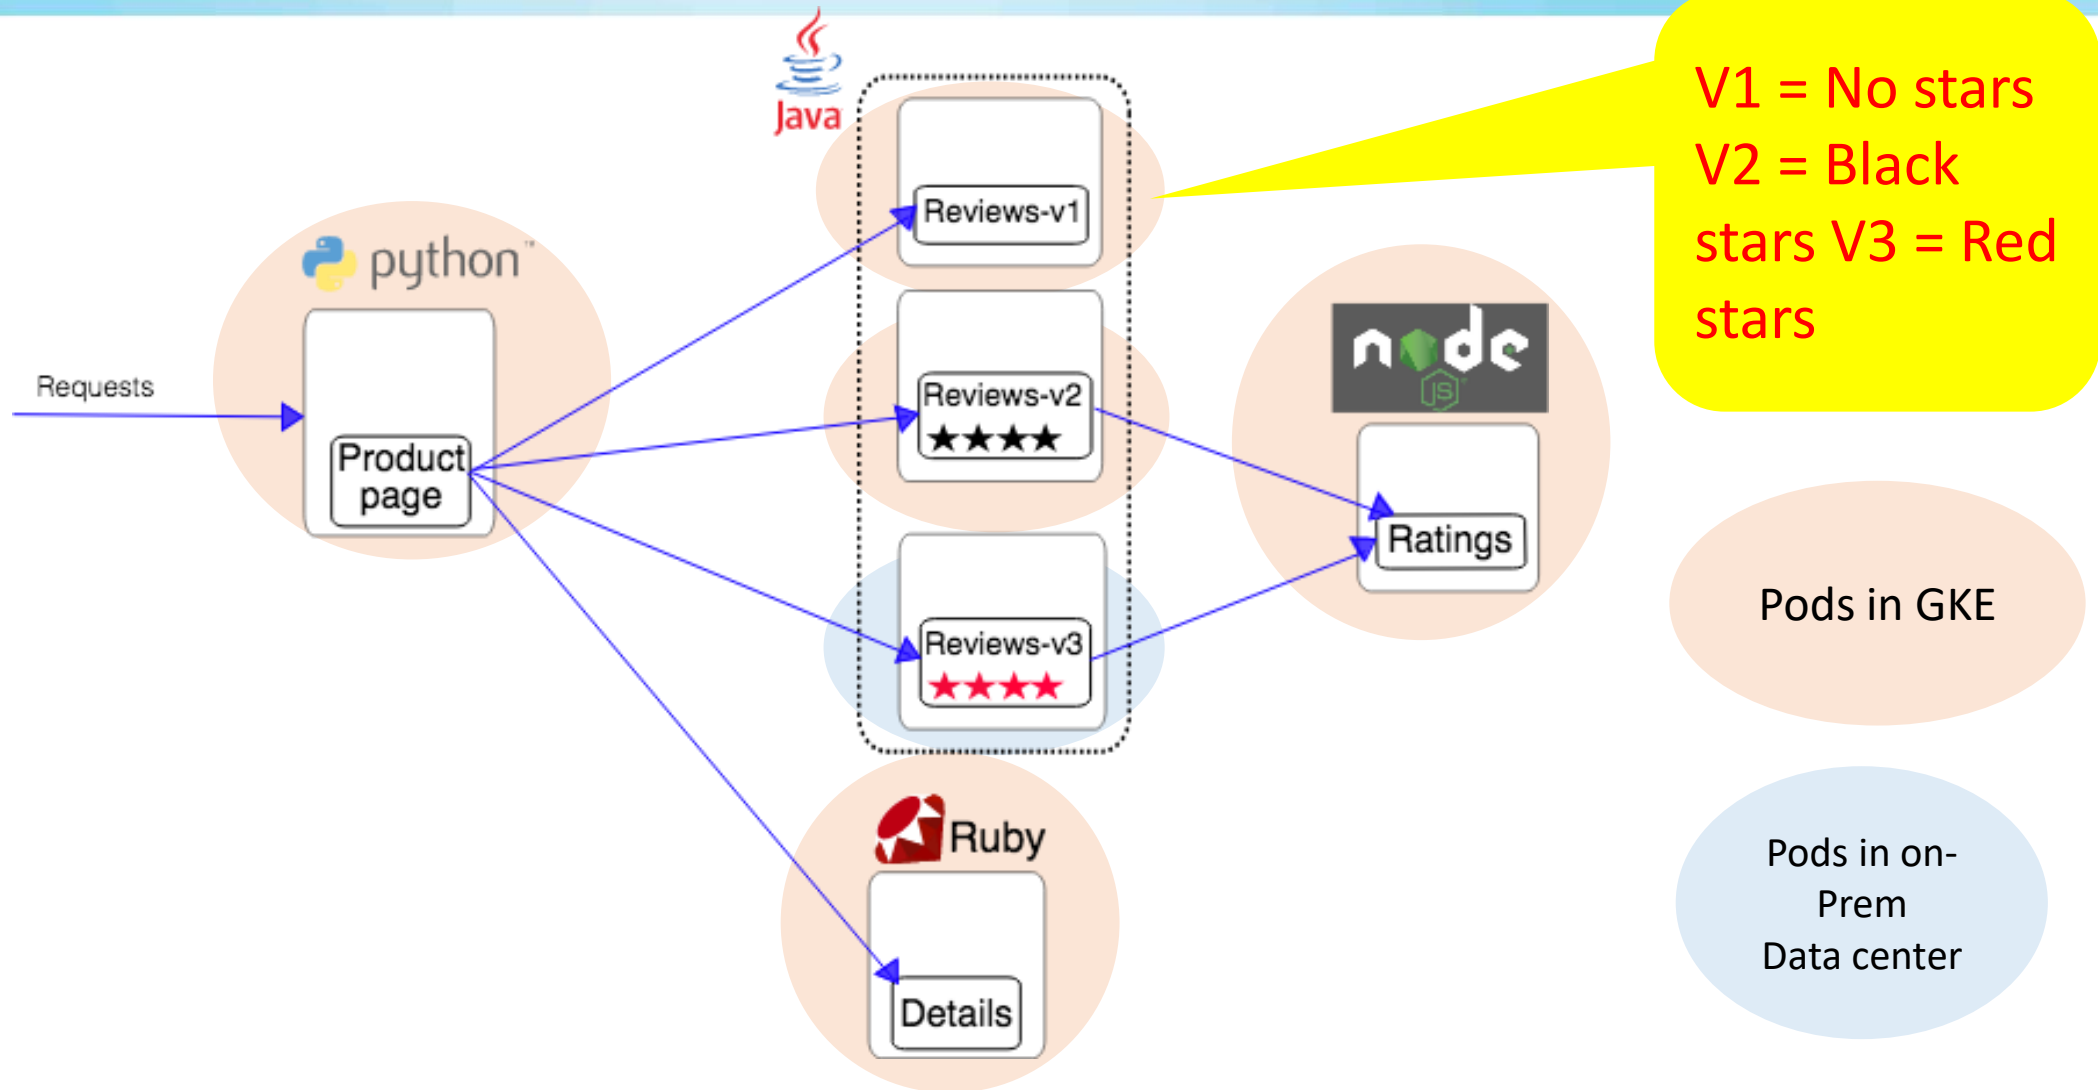


KubeCon



CloudNativeCon

Europe 2018






# Bookinfo Application Flow






Europe 2018

End User Request  
For Product page

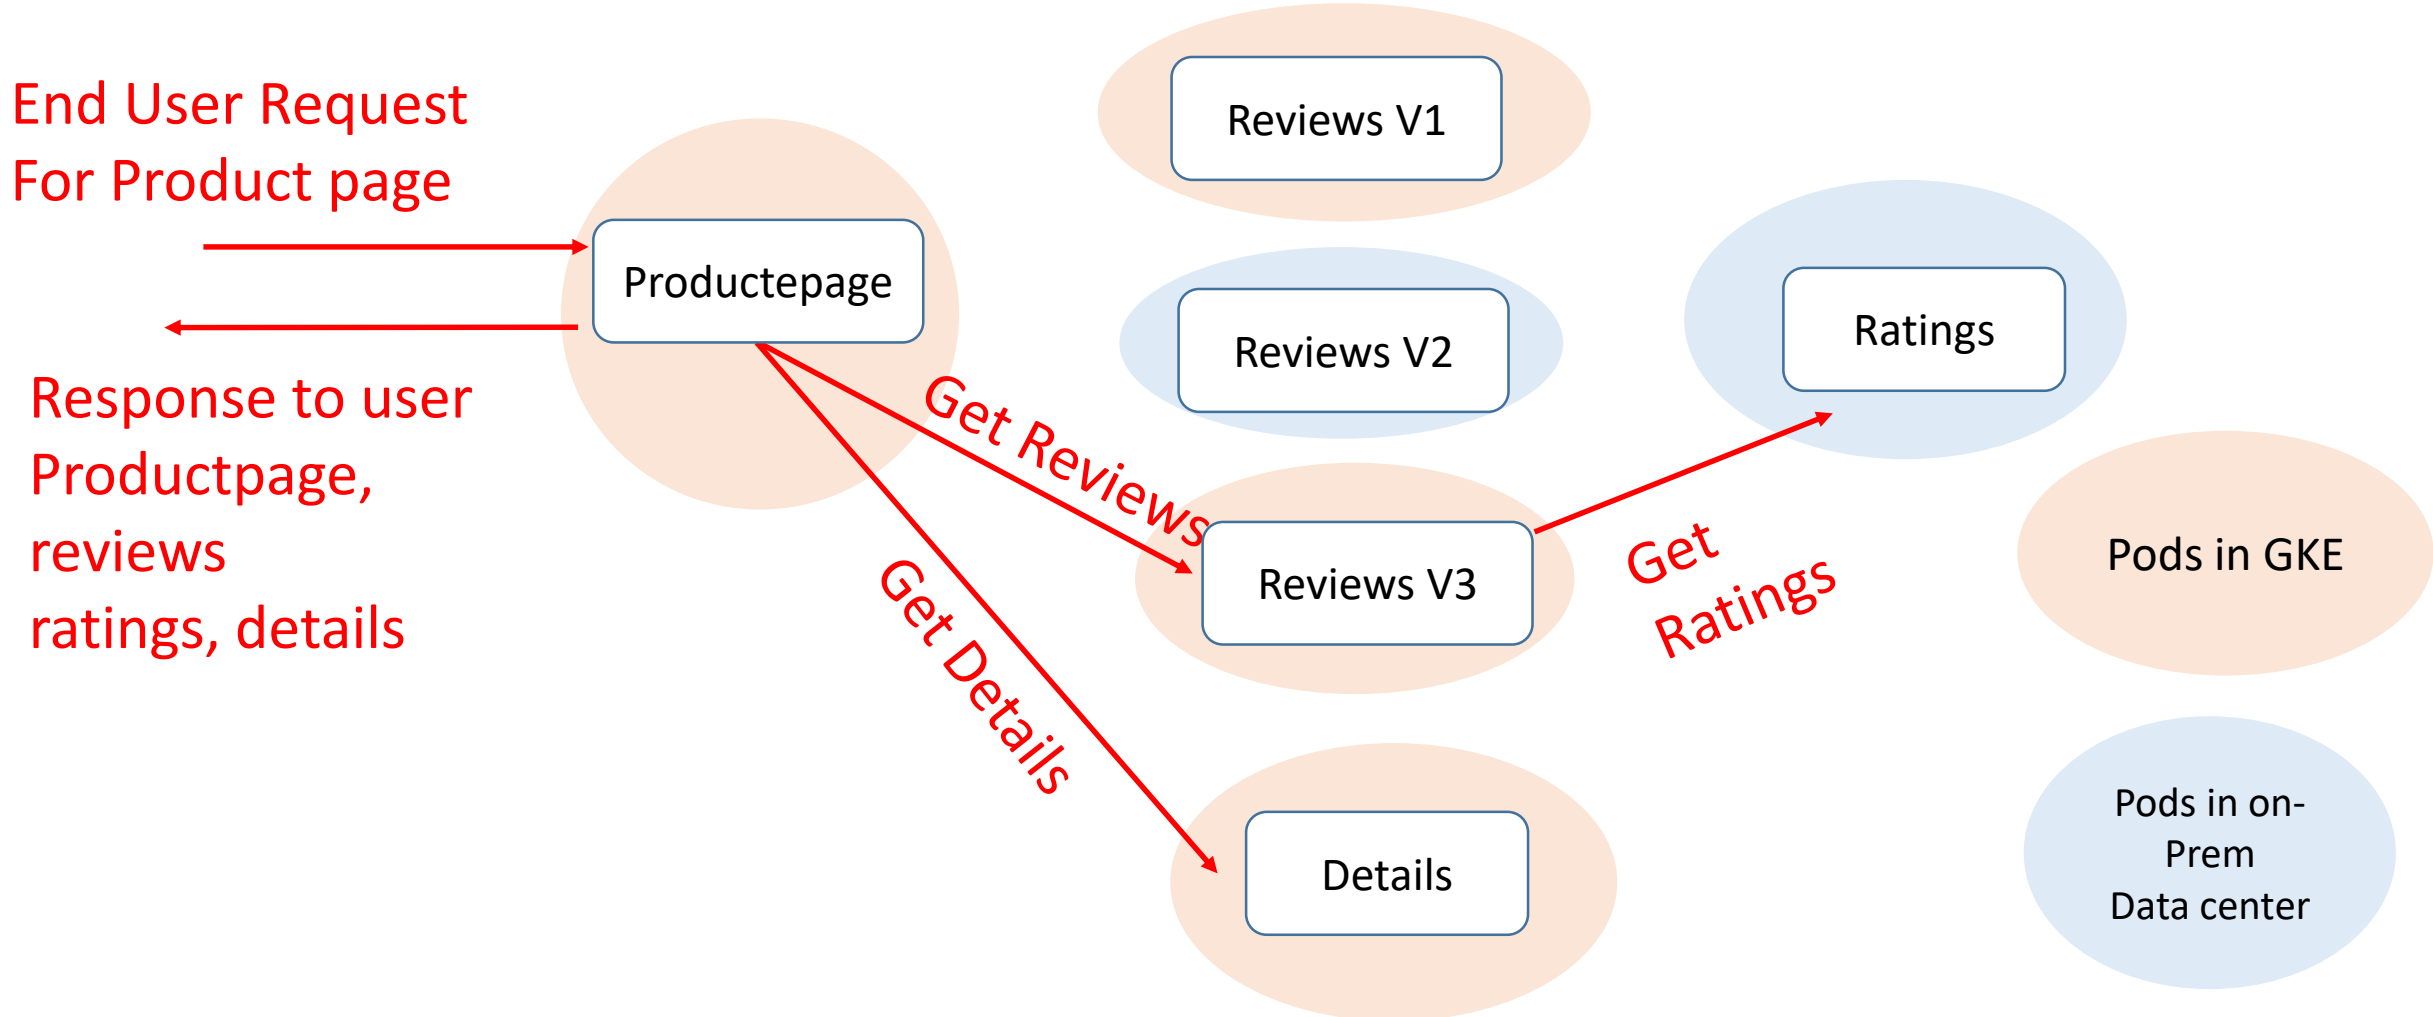


Response to user  
Productpage,  
reviews  
ratings, details



Get Reviews  
Get Details

Get Ratings



Deployment in action

```
GKE # kubectl get namespaces
NAME          STATUS    AGE
default       Active    3h
kube-public   Active    3h
kube-system   Active    3h
GKE #
```

```
ON-PREM k8s # kubectl get namespaces
NAME          STATUS    AGE
default       Active    5h
kube-public   Active    5h
kube-system   Active    5h
ON-PREM k8s #
```

```
Terminal Shell Edit View Window Help
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x19
GKE # kubectl get namespaces
NAME          STATUS    AGE
default       Active   3h
kube-public   Active   3h
kube-system   Active   3h
GKE # make e2e_bookinfo E2E_ARGS="--skip_cleanup --cluster_registry_dir=/root/kubeconf --namespace=istio-system" | tee $HOME/pilot_log
```

```
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18
ON-PREM k8s # kubectl get namespaces
NAME          STATUS    AGE
default       Active   5h
kube-public   Active   5h
kube-system   Active   5h
ON-PREM k8s #
```

```
Terminal Shell Edit View Window Help
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x19
kube-system Active 3h
GKE # make e2e_bookinfo E2E_ARGS="--skip_cleanup --cluster_registry_dir=/root/kubeconf --namespace=istio-system" | tee $HOME/pilot_log
bin/gobuild.sh /root/go/out/linux_amd64/release/istioctl istio.io/istio/pkg/version ./istioctl/cmd/istioctl

real    0m0.412s
user    0m0.548s
sys     0m0.156s
./install/updateVersion.sh -a docker.io/johnajoyce,kubeconfig_fix
/tmp/templates ~/go/src/istio.io/istio
~/go/src/istio.io/istio
/tmp/templates/addons ~/go/src/istio.io/istio
~/go/src/istio.io/istio
/tmp/templates ~/go/src/istio.io/istio
~/go/src/istio.io/istio
/tmp/templates ~/go/src/istio.io/istio
~/go/src/istio.io/istio
-a docker.io/johnajoyce,kubeconfig_fix
make[1]: Entering directory '/root/go/src/istio.io/istio'
```

```
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18
ON-PREM k8s # kubectl get namespaces
NAME          STATUS   AGE
default       Active   5h
kube-public   Active   5h
kube-system   Active   5h
ON-PREM k8s #
```

```
Terminal Shell Edit View Window Help
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x19
2018-04-25T19:51:46.064141Z      info    Running command kubectl logs reviews-v3-6cb5984c75-hfxk8 -n istio-system -c istio-proxy --kubeconfig=/root/kubeconf/local
2018-04-25T19:51:46.341756Z      info    Running command kubectl logs reviews-v3-6cb5984c75-hfxk8 -n istio-system -c istio-proxy -p --kubeconfig=/root/kubeconf/local
2018-04-25T19:51:46.589116Z      info    Command error: exit status 1
2018-04-25T19:51:46.589164Z      info    No previous log command failed: "Error from server (BadRequest): previous terminated container \"istio-proxy\" in pod \"reviews-v3-6cb5984c75-hfxk8\" not found\n" exit status 1
2018-04-25T19:51:46.589183Z      info    Fetching deployment info on pod

2018-04-25T19:51:46.589197Z      info    Running command kubectl get pod -n istio-system -o yaml --kubeconfig=/root/kubeconf/local
2018-04-25T19:51:46.907629Z      info    Fetching deployment info on service

2018-04-25T19:51:46.907661Z      info    Running command kubectl get service -n istio-system -o yaml --kubeconfig=/root/kubeconf/local
2018-04-25T19:51:47.120189Z      info    Fetching deployment info on ingress

2018-04-25T19:51:47.120227Z      info    Running command kubectl get ingress -n istio-system -o yaml --kubeconfig=/root/kubeconf/local
2018-04-25T19:51:47.292585Z      info    Dev mode (--skip_cleanup), skipping cleanup (removal of namespace/install)
ok      istio.io/istio/tests/e2e/tests/bookinfo 187.069s
GKE #
```

```
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18
ON-PREM k8s # kubectl get namespaces
NAME          STATUS   AGE
default       Active   5h
kube-public   Active   5h
kube-system   Active   5h
ON-PREM k8s #
```

```
Terminal Shell Edit View Window Help
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x19

istio-citadel-c8dbbc878-5drxn          1/1      Running   0         4m
istio-egressgateway-789d849cb-97sg4    1/1      Running   0         7m
istio-ingress-6b8784d5df-5ztj8         1/1      Running   0         7m
istio-ingressgateway-68b9c87b66-fczzp  1/1      Running   0         7m
istio-mixer-create-cr-rmkrt           0/1      Completed 0         7m
istio-pilot-69554969b8-89n8p          2/2      Running   0         7m
istio-policy-659b5c984f-lzmgc         2/2      Running   0         7m
istio-statsd-prom-bridge-6dbb7dcc7f-2qh4b 1/1      Running   0         7m
istio-telemetry-57944d9644-87ffs      2/2      Running   0         7m
mongodb-v1-9d4cd8bb9-ml7zz           2/2      Running   0         6m
mysqldb-v1-7f8bfbbf6c-wtx5s          2/2      Running   0         6m
productpage-v1-597789bc47-ssmls       2/2      Running   0         6m
prometheus-586d95b8d9-rts55          1/1      Running   0         7m
ratings-v1-6cdf5c76f5-6qvcs          2/2      Running   0         6m
ratings-v2-b67487df8-rn2lq           2/2      Running   0         6m
ratings-v2-mysql-59ddf6fbd-cspbz      2/2      Running   0         6m
reviews-v1-d8877bf75-2q2sr           2/2      Running   0         6m
reviews-v2-64d684d755-7flp6          2/2      Running   0         6m
GKE #
```

```
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18

ON-PREM k8s # kubectl get pods -n istio-system
NAME                                READY    STATUS    RESTARTS   AGE
istio-citadel-599f4bb46-qnh66      1/1     Running   0           1m
ON-PREM k8s #
```

```
GKE # kubectl get services -n istio-system
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
details	ClusterIP	10.125.35.188	<none>	9080/TCP	3m
istio-citadel	ClusterIP	10.125.46.223	<none>	8060/TCP,9093/TCP	4m
istio-egressgateway	ClusterIP	10.125.40.2	<none>	80/TCP,443/TCP	4m
istio-ingress	LoadBalancer	10.125.60.5	35.203.186.136	80:30705/TCP,443:31891/TCP	4m
istio-ingressgateway	NodePort	10.125.22.148	<none>	80:31380/TCP,443:31390/TCP,31400:31400/TCP	4m
istio-pilot	ClusterIP	10.125.4.161	<none>	15003/TCP,15005/TCP,15007/TCP,15010/TCP,15011/TCP,8080/TCP,9093/TCP	4m
istio-policy	ClusterIP	10.125.35.206	<none>	9091/TCP,15004/TCP,9093/TCP	4m
istio-statsd-prom-bridge	ClusterIP	10.125.52.238	<none>	9102/TCP,9125/UDP	4m
istio-telemetry	ClusterIP	10.125.43.186	<none>	9091/TCP,15004/TCP,9093/TCP,42422/TCP	4m
mongodb	ClusterIP	10.125.59.136	<none>	27017/TCP	3m
mysqldb	ClusterIP	10.125.36.63	<none>	3306/TCP	3m
productpage	ClusterIP	10.125.9.95	<none>	9080/TCP	3m
prometheus	ClusterIP	10.125.55.96	<none>	9090/TCP	4m
ratings	ClusterIP	10.125.33.214	<none>	9080/TCP	3m
reviews	ClusterIP	10.125.59.1	<none>	9080/TCP	3m

```
GKE #
```

```
ON-PREM k8s # kubectl get services -n istio-system
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
istio-citadel	ClusterIP	10.101.29.255	<none>	8060/TCP,9093/TCP	4m
istio-pilot	ClusterIP	None	<none>	15003/TCP,15005/TCP,15007/TCP,15010/TCP,15011/TCP,8080/TCP,9093/TCP	4m
istio-policy	ClusterIP	None	<none>	9091/TCP,15004/TCP,9093/TCP,9094/TCP,9102/TCP,9125/UDP,42422/TCP	4m
istio-statsd-prom-bridge	ClusterIP	None	<none>	9102/TCP,9125/UDP	4m

```
ON-PREM k8s #
```



```

reviews                ClusterIP      10.125.59.1    <none>          9080/TCP
GKE # kubectl get endpoints -n istio-system
NAME                   ENDPOINTS
details                10.124.0.18:9080,10.124.2.32:9080
istio-citadel          10.124.2.33:9093,10.124.2.33:8060
istio-egressgateway   10.124.2.22:80,10.124.2.22:443
istio-ingress          10.124.2.23:80,10.124.2.23:443
istio-ingressgateway  10.124.2.24:80,10.124.2.24:31400,10.124.2.24:443
istio-pilot            10.124.2.26:9093,10.124.2.26:15010,10.124.2.26:15007 + 4 more...
istio-policy           10.124.0.16:9093,10.124.0.16:9091,10.124.0.16:15004
istio-statsd-prom-bridge 10.124.2.21:9125,10.124.2.21:9102
istio-telemetry        10.124.1.15:42422,10.124.1.15:9093,10.124.1.15:9091 + 1 more...
mongodb                10.124.2.31:27017
mysqldb                10.124.0.21:3306
productpage            10.124.2.29:9080
prometheus             10.124.0.17:9090
ratings                10.124.0.20:9080,10.124.1.17:9080,10.124.2.30:9080
reviews                10.124.0.19:9080,10.124.1.18:9080,10.124.2.28:9080
GKE #

```

```

ON-PREM k8s # kubectl get services -n istio-system
NAME                   TYPE          CLUSTER-IP      EXTERNAL-IP      PORT(S)
istio-citadel          ClusterIP     10.101.29.255   <none>            8060/TCP,9093/TCP
istio-pilot            ClusterIP     None             <none>            15003/TCP,15005/TCP,15007/TCP,15010/TCP,15011/TCP,8080/TCP,9093/TCP
istio-policy           ClusterIP     None             <none>            9091/TCP,15004/TCP,9093/TCP,9094/TCP,9102/TCP,9125/UDP,42422/TCP
istio-statsd-prom-bridge ClusterIP     None             <none>            9102/TCP,9125/UDP
ON-PREM k8s # kubectl get endpoints -n istio-system
NAME                   ENDPOINTS
istio-citadel          192.170.228.83:9093,192.170.228.83:8060
istio-pilot            10.124.2.26:9093,10.124.2.26:15010,10.124.2.26:15007 + 4 more...
istio-policy           10.124.0.16:42422,10.124.0.16:9091,10.124.0.16:9093 + 4 more...
istio-statsd-prom-bridge 10.124.2.21:9125,10.124.2.21:9102
ON-PREM k8s #

```

```
GKE # kubectl get services -n istio-system
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
details	ClusterIP	10.125.35.188	<none>	9080/TCP	17m
istio-citadel	ClusterIP	10.125.46.223	<none>	8060/TCP,9093/TCP	19m
istio-egressgateway	ClusterIP	10.125.40.2	<none>	80/TCP,443/TCP	19m
istio-ingress	LoadBalancer	10.125.60.5	35.203.186.136	80:30705/TCP,443:31891/TCP	19m
istio-ingressgateway	NodePort	10.125.22.148	<none>	80:31380/TCP,443:31390/TCP,31400:31400/TCP	19m
istio-pilot	ClusterIP	10.125.4.161	<none>	15003/TCP,15005/TCP,15007/TCP,15010/TCP,15011/TCP,8080/TCP,9093/TCP	19m
istio-policy	ClusterIP	10.125.35.206	<none>	9091/TCP,15004/TCP,9093/TCP	19m
istio-statsd-prom-bridge	ClusterIP	10.125.52.238	<none>	9102/TCP,9125/UDP	19m
istio-telemetry	ClusterIP	10.125.43.186	<none>	9091/TCP,15004/TCP,9093/TCP,42422/TCP	19m
mongodb	ClusterIP	10.125.59.136	<none>	27017/TCP	17m
mysqldb	ClusterIP	10.125.36.63	<none>	3306/TCP	17m
productpage	ClusterIP	10.125.9.95	<none>	9080/TCP	17m
prometheus	ClusterIP	10.125.55.96	<none>	9090/TCP	19m
ratings	ClusterIP	10.125.33.214	<none>	9080/TCP	17m
reviews	ClusterIP	10.125.59.1	<none>	9080/TCP	17m

```
GKE #
```

```
ON-PREM k8s #
```

```
Terminal Shell Edit View Window Help
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x19
GKE #
```

```
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18
ON-PREM k8s # kubectl apply -f ratings-svc.yaml -n istio-system
service "ratings" created
ON-PREM k8s # kubectl apply -f reviews-v3_inject.yaml -n istio-system
service "reviews" created
deployment.extensions "reviews-v3" created
ON-PREM k8s #
```

```
Terminal Shell Edit View Window Help
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x19
GKE #
```

```
tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18
ON-PREM k8s # kubectl get pods -n istio-system
NAME                READY   STATUS             RESTARTS   AGE
istio-citadel-599f4bb46-qnh66  1/1    Running            0           1h
reviews-v3-7d9bc85dcb-5468r    0/2    PodInitializing    0           39s
ON-PREM k8s # kubectl get pods -n istio-system
NAME                READY   STATUS    RESTARTS   AGE
istio-citadel-599f4bb46-qnh66  1/1    Running   0           1h
reviews-v3-7d9bc85dcb-5468r    2/2    Running   0           55s
ON-PREM k8s # kubectl get svc -n istio-system
NAME                TYPE        CLUSTER-IP    EXTERNAL-IP    PORT(S)                AGE
istio-citadel       ClusterIP   10.101.29.255 <none>         8060/TCP,9093/TCP      1h
istio-pilot         ClusterIP   None          <none>         15003/TCP,15005/TCP,15007/TCP,15010/TCP,15011/TCP,8080/TCP,9093/TCP 1h
istio-policy        ClusterIP   None          <none>         9091/TCP,15004/TCP,9093/TCP,9094/TCP,9102/TCP,9125/UDP,42422/TCP 1h
istio-statsd-prom-bridge ClusterIP   None          <none>         9102/TCP,9125/UDP      1h
ratings             ClusterIP   10.110.96.10  <none>         9080/TCP                1m
reviews             ClusterIP   10.111.253.54 <none>         9080/TCP                1m
ON-PREM k8s #
```

```
GKE # PILOT_URL=10.124.2.26:8080 ~/istio-proxy-cfg c s productpage | grep -A 10 -B 1 reviews
```

```
% Total      % Received % Xferd  Average Speed   Time    Time       Time  Current
                               Dload  Upload  Total   Spent    Left   Speed
100 10423      0 10423      0      0   193k      0  --:--:--  --:--:--  --:--:--  195k
```

```
{
  "name": "out.reviews.istio-system.svc.cluster.local|http",
  "service_name": "reviews.istio-system.svc.cluster.local|http",
  "connect_timeout_ms": 1000,
  "type": "sds",
  "lb_type": "round_robin"
},
{
  "name": "out.sleep.default.svc.cluster.local|http",
  "service_name": "sleep.default.svc.cluster.local|http",
  "connect_timeout_ms": 1000,
  "type": "sds",
  "lb_type": "round_robin"
}
```

```
GKE # █
```

tiswanso — root@k8s-demo: ~ — ssh · ssh root@k8s-demo — 157x18

```
ON-PREM k8s # kubectl get pods -n istio-system -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
istio-citadel-599f4bb46-qnh66	1/1	Running	0	1h	192.170.228.83	k8s-demo
reviews-v3-7d9bc85dcb-5468r	2/2	Running	0	22m	192.170.228.84	k8s-demo

```
ON-PREM k8s # █
```

```
port": 9080,  
  "tags": {  
    "az": "us-west1/us-west1-a"  
  }  
},  
{  
  "ip_address": "10.124.2.34",  
  "port": 9080,  
  "tags": {  
    "az": "us-west1/us-west1-a"  
  }  
},  
{  
  "ip_address": "192.170.228.84",  
  "port": 9080  
}  
]  
}  
GKE # █
```

tiswanso — root@k8s-demo: ~ — ssh • ssh root@k8s-demo — 157x18

```
ON-PREM k8s # kubectl get pods -n istio-system -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
istio-citadel-599f4bb46-qnh66	1/1	Running	0	1h	192.170.228.83	k8s-demo
reviews-v3-7d9bc85dcb-5468r	2/2	Running	0	22m	192.170.228.84	k8s-demo

# Mesh Routing Demo





Future Improvements

# Future improvements:

- Usability Improvements
- Increased test coverage
- Pilot to Pilot multi-cloud
- Full adoption of the clusterregistry (via API calls from Istio)
- Zero VPN work
- Bring your own Service Name
- Mixer Multicluster enabled

Q & A

# References:



KubeCon



CloudNativeCon

Europe 2018

- Initial PR to enable Multicluster: <https://github.com/istio/istio/pull/2880>
- Installation documentation: <https://github.com/istio/istio.github.io/pull/1139>
- Multicloud Design Document: <https://tinyurl.com/y7scozvb>
- Zero VPN design document: <https://tinyurl.com/zerovpn>
- Bring Your own Service Name Design Document: <https://tinyurl.com/svc-name>