

What is Istio ?



Abdellfetah SGHIOUAR

Strategic Cloud Engineer @Google Stockholm Twitter: boredabdel@

Today we're going to cover the why, what, and how of lstio



Why Istio ?





Containers?

Great!





Kubernetes?

Also Great!





But microservices deployments are hard



How do you go from a bunch of services...



...to a well organized and functioning deployment?



What makes microservices difficult?





What is Istio?





How does Istio help?







Uniform observability Operational agility

Policy driven security



How does Istio help?

Telemetry	Traffic	Security
Examine everything	Manage the flow of	Secure access and
happening with your	traffic into, out of ,	communications
services with little to	and within your	between some or all
no instrumentation	complex deployments	services



lstio at a glance





lstio sidecar proxy





lstio data plane





lstio control plane





What does Istio do?





Observing services

Get automatic

tracing, monitoring,

and logging of all your

services.





Connecting services

Using VirtualService, DestinationRule, Gateway, and ServiceEntry objects, Istio helps you with:







independent of number of instances supporting the version



Content-based traffic steering - The content of a request can be used to determine the destination of a request



Connecting services

Fault injection and

circuit breaking





Securing services

Automatically secure your services through managed authentication, authorization, and encryption of communication between services.





Controlling and securing services

Apply broad or fine-grained security policies to some or all of your workloads



Istio Security Architecture



How do you put Istio to work?





Including the sidecar proxy in your Pods

manually injecting the sidecar proxy kubectl apply -f <(istioctl kube-inject -f deployment.yaml)</pre>

auto-injecting the sidecar proxy kubectl label ns default istio-injection=enabled



Traffic management





Simple deployment

This is a simple

frontend/backend

deployment. We want to canary

test v2 of the backend. Each

backend provides slightly

different functionality.





Recap: Traffic splitting



Traffic splitting decoupled from infrastructure scaling - proportion of traffic routed to a version is independent of number of instances supporting the version



Distributing traffic

Without sidecar proxies, how can we use Istio's traffic distribution? The frontend should hit v1 90% of the time, and v2 10% of the time.





Default Kubernetes round-robin routing





Traffic policy for backend deployment

Create a DestinationRule policy for weather-backend - in this case, creating named subsets to send traffic to apiVersion: networking.istio.io/v1alpha3 kind: DestinationRule metadata: name: weather-backend-destination spec: host: weather-backend subsets: - name: single labels: version: single - name: multiple labels: version: multiple



Routing traffic from ingress

Create routing rules using a VirtualService, telling the Gateway object where/how to distribute traffic



apiVersion: networking.istio.io/v1alpha3
kind: VirtualService
metadata:

name: weather-backend-service
spec:

hosts:

- "*"

gateways:

- weather-backend-gateway
http:

- match:

- uri:

exact: /api/weather
route:

- destination: host: weather-backend subset: single port: number: 5000 weight: 90 - destination: host: weather-backend subset: multiple port: number: 5000

weight: 10

Securing services





Securing a subset of services

How do you use Istio to slowly deploy mTLS across services, while also keeping legacy clients in mind?





Rolling out mTLS

- 1. Deploy Istio with PERMISSIVE mTLS settings
- 2. Apply Policy for Service A with PERMISSIVE mode
- 3. Apply DestinationRule for Service A with MUTUAL mode
 - a. Services **with** istio-proxy encrypt traffic using mTLS
 - b. Services **without** istio-proxy don't encrypt traffic
- 4. When ready, apply Policy for Service A with STRICT mode



Apply Policy with PERMISSIVE mode

apiVersion: auth.istio.io/v1alpha1
kind: Policy
metadata:
 name: mtls-backend
 namespace: secure
spec:
 targets:
 - name: weather-backend
 peers:
 - mtls:

mode: PERMISSIVE





Apply DestinationRule with MUTUAL mode

```
apiVersion: net.istio.io/v1alpha3
kind: DestinationRule
metadata:
   name: mtls-mutual
spec:
   host: weather-backend.secure
   trafficPolicy:
     tls:
        mode: ISTIO_MUTUAL
```





Apply Policy with STRICT mode

apiVersion: auth.istio.io/v1alpha1
kind: Policy
metadata:
 name: mtls-backend
 namespace: secure
spec:
 targets:
 - name: weather-backend
 peers:
 - mtls:

mode: STRICT





Gathering telemetry





Telemetry for free

Without any instrumentation, Istio captures a predefined set of **metrics**, request **traces**, and service **logs** - and forwards them to the configured adapters.



Istio built-in metrics

Request Count Incremented for every request handled

Request Duration Duration of the request

Request Size Size of the HTTP request body **Response Size** Size of the HTTP response body

TCP Byte Sent Total bytes sent during response

TCP Byte Received Total bytes received during request



Where do you go from here ?





Some light reading to get you started

Service Mesh Era blog posts (conceptual and practical intros)

Incremental Istio: Traffic Management

mTLS Migration

Istio Concepts

github.com/crscmnky/weatherinfo

github.com/crcsmnky/next19-incremental-istio



Thank you!

