

PagerDuty

Service Messy

Pagerduty's Service Mesh Journey

Liz Frost, SRE IV

PagerDuty Core Infrastructure

Liz Frost, SRE IV

PagerDuty Core Infrastructure



On pronunciation

/kʊbɜː'nɛtɪz/ or /kʊbɜː'nɛtɪz

VS

kə'bɜːnɛtɪz

- **What to know before you go**
- **How we migrated**
- **Random Gotchas**

What We Have

- Consul
- Kubernetes (from Nomad)
- Lots of Elixir
- Lots of *other* languages
- An “Opinionated Wrapper Library”
- A lot of small rocks and some *really big ones*

Eventual Consistency

App Types	Total per-app Type to migrate
Elixir	100
Shell	10
Python	23
Envoy	0
Nginx	4
Ruby	6
Scala	7
Go	1
Java	5
TypeScript	4
JavaScript	2
C#	3
N/A	0

Problems:

- Too many languages
- Too many implementations
- Compliance issues

PagerDuty

Mesh Time

(know before you go)

Pros

- Increased Visibility
- Security
- Traffic Features

Cons

- Complexity
- Latency
- Resource Intensive
- Learning Curve

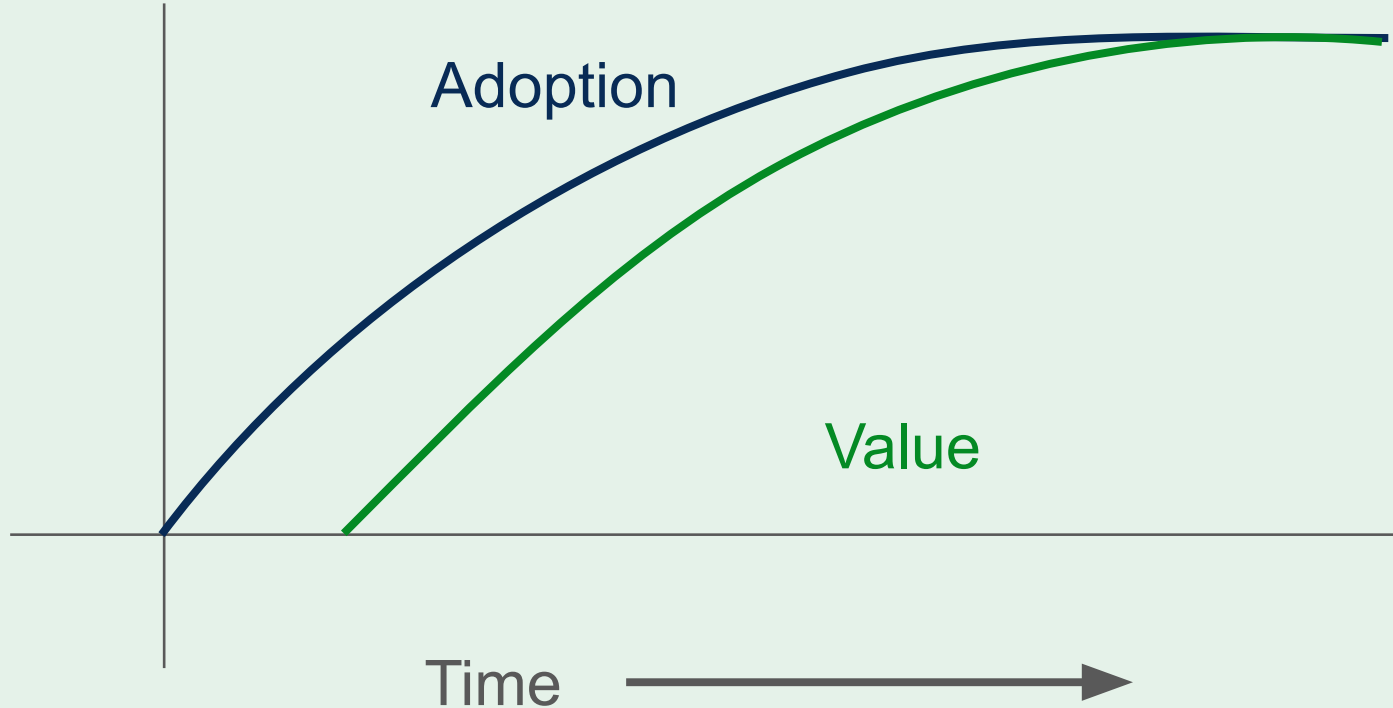
Pros

- Increased Visibility
- Security
- Traffic Features
 - Retries, Canaries, Circuit breakers...

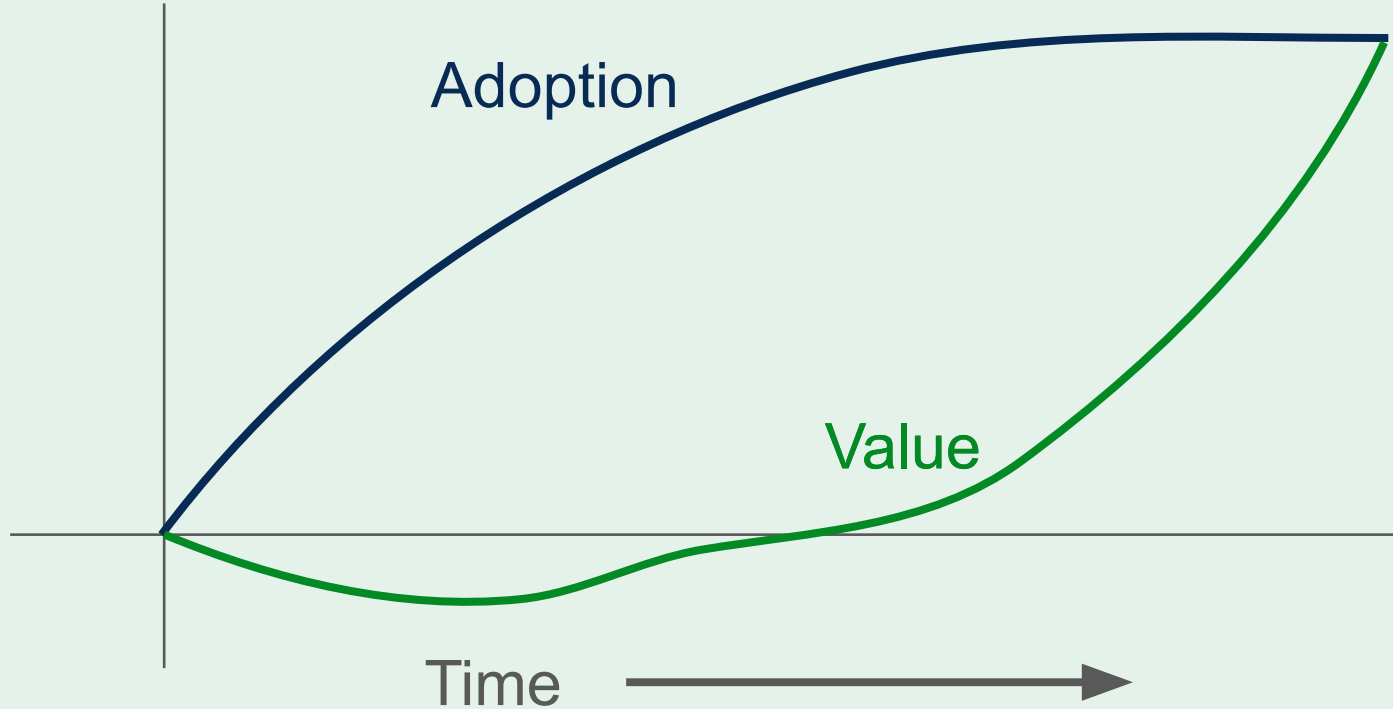
Cons

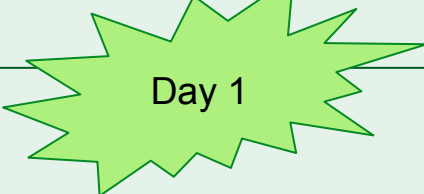
- Complexity
- Latency
- Resource Intensive
- Learning Curve

Network Effect :)



Network Effect :(





Day 1

Pros

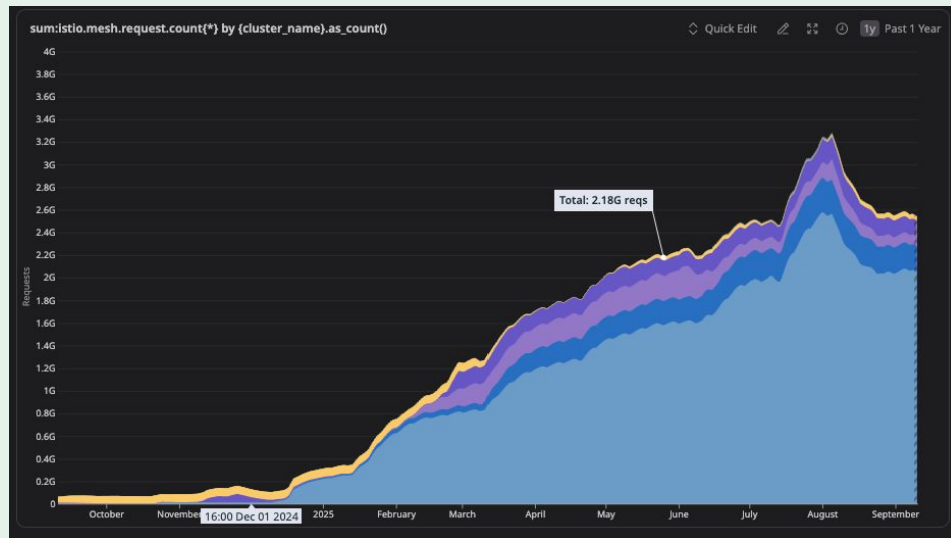
- Increased Visibility
- ~~● Security~~
- ~~● Traffic Features~~

Cons

- Complexity
- Latency
- Resource Intensive
- Learning Curve

Metrics work

- Source or destination
- Normalised across every app
 - No guessing metric names!
- Graph go up



Security

- No Enforcement
- Partial Encryption?
more like no encryption :(
- No Validation

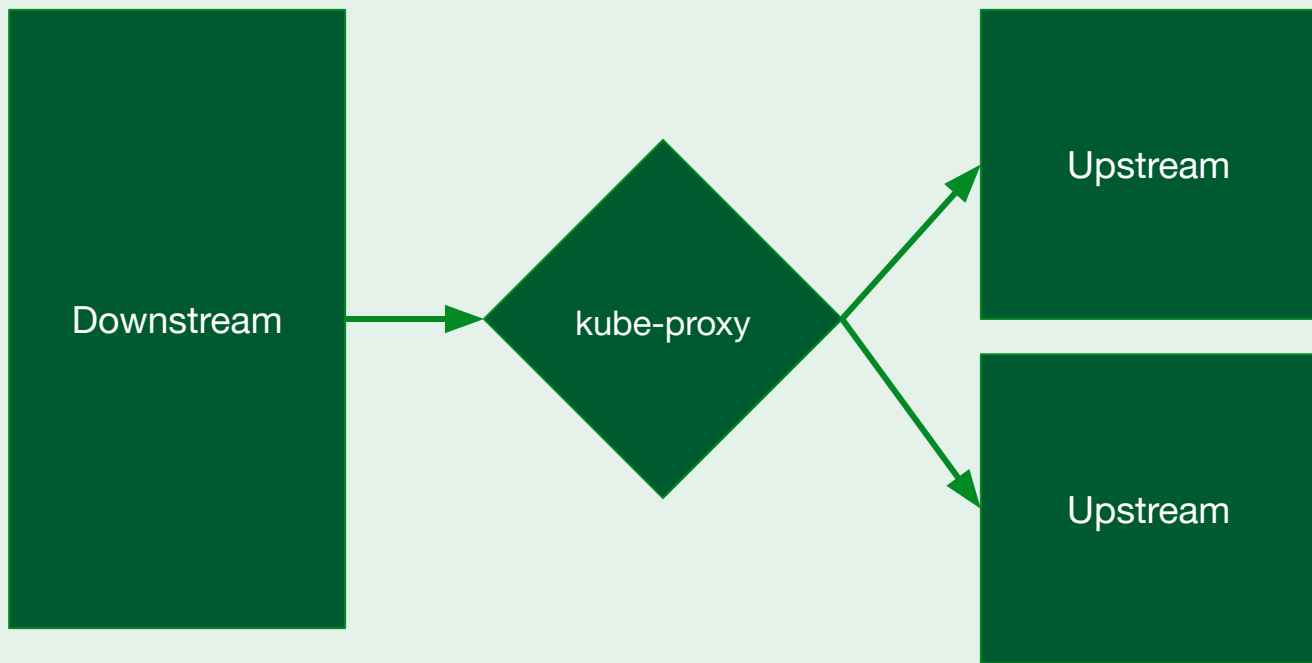
ssl added and removed
everywhere :(

Traffic Features?

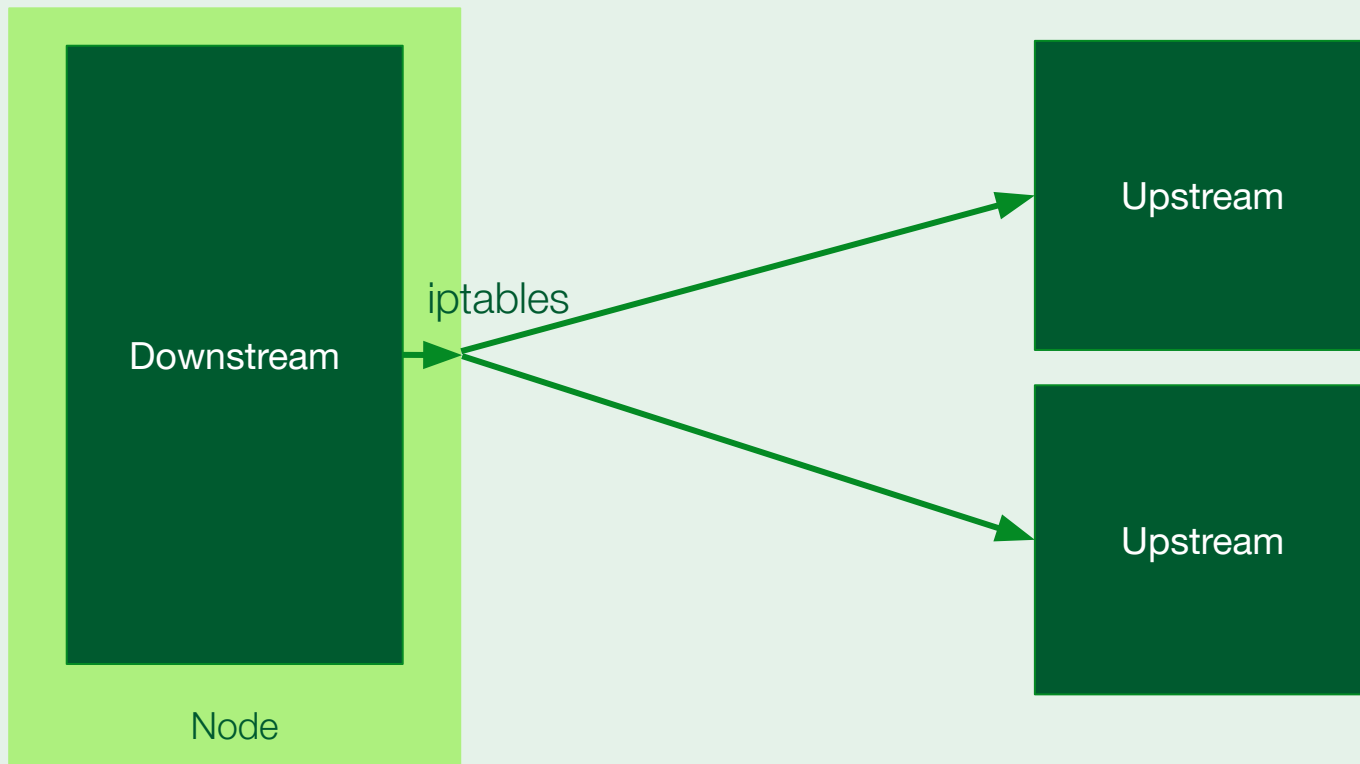
A Brief Detour...

```
apiVersion: v1
kind: Service
metadata:
  name: pony-api
  namespace: pony-api
spec:
  ports:
  - name: http
    port: 80
    protocol: TCP
    targetPort: 8080
  selector:
    app: pony-api
  type: ClusterIP
```

How you think about it



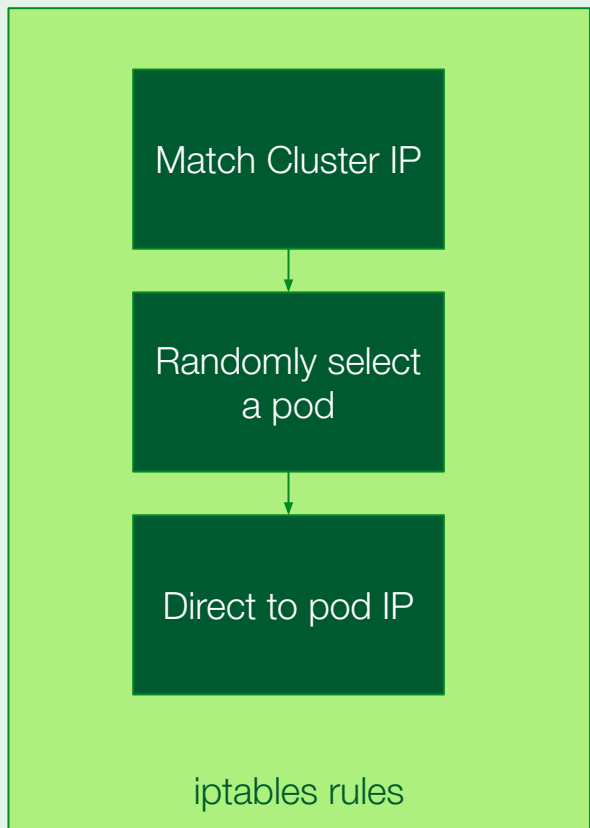
How you think about it



the cluster IP does not exist!

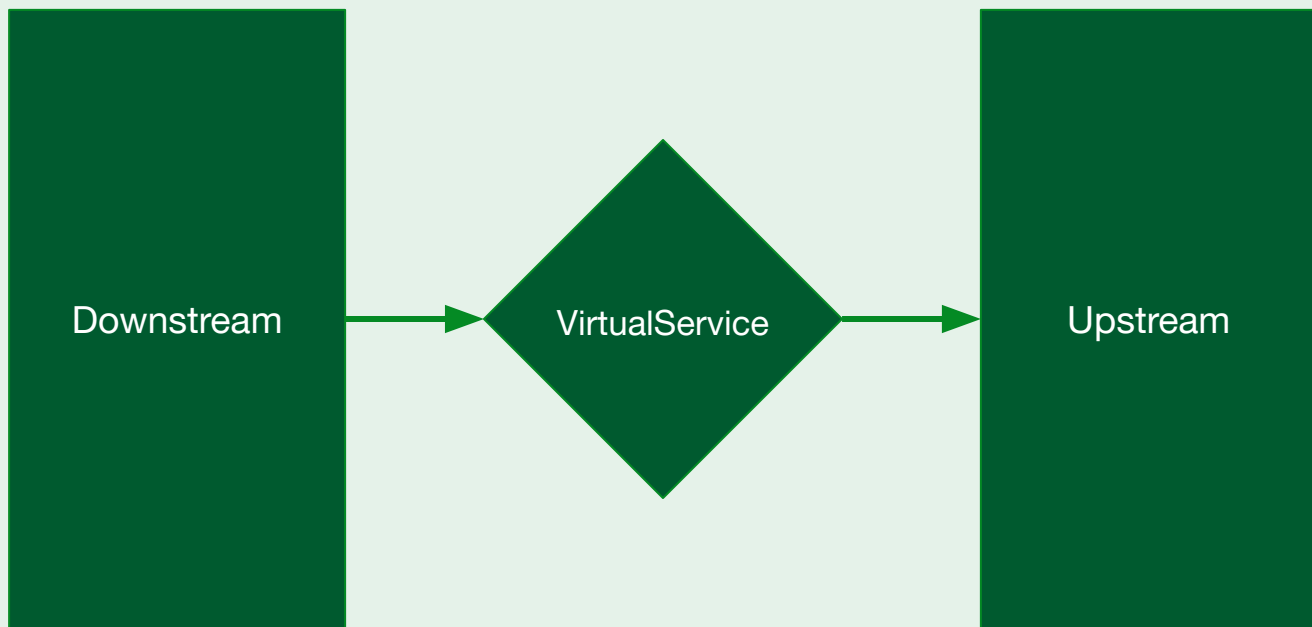
kube-“proxy”

it's all iptables

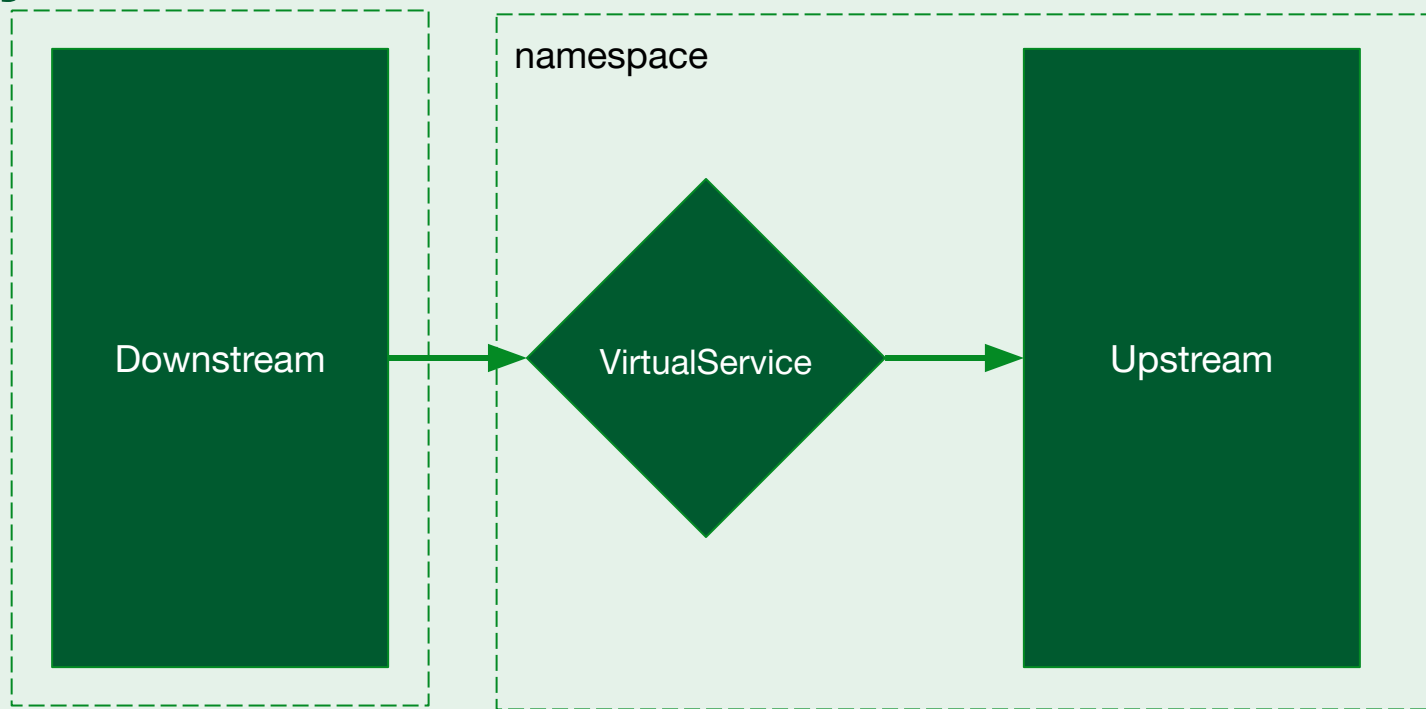


```
apiVersion: networking.istio.io/v1
kind: VirtualService
metadata:
  name: pony-api
spec:
  hosts:
  - pony-api.pony-api.svc.cluster.local
  http:
  - route:
    - destination:
        host: pony-api.pony-api.svc.cluster.local
        port:
          number: 80
    retries:
      attempts: 3
```

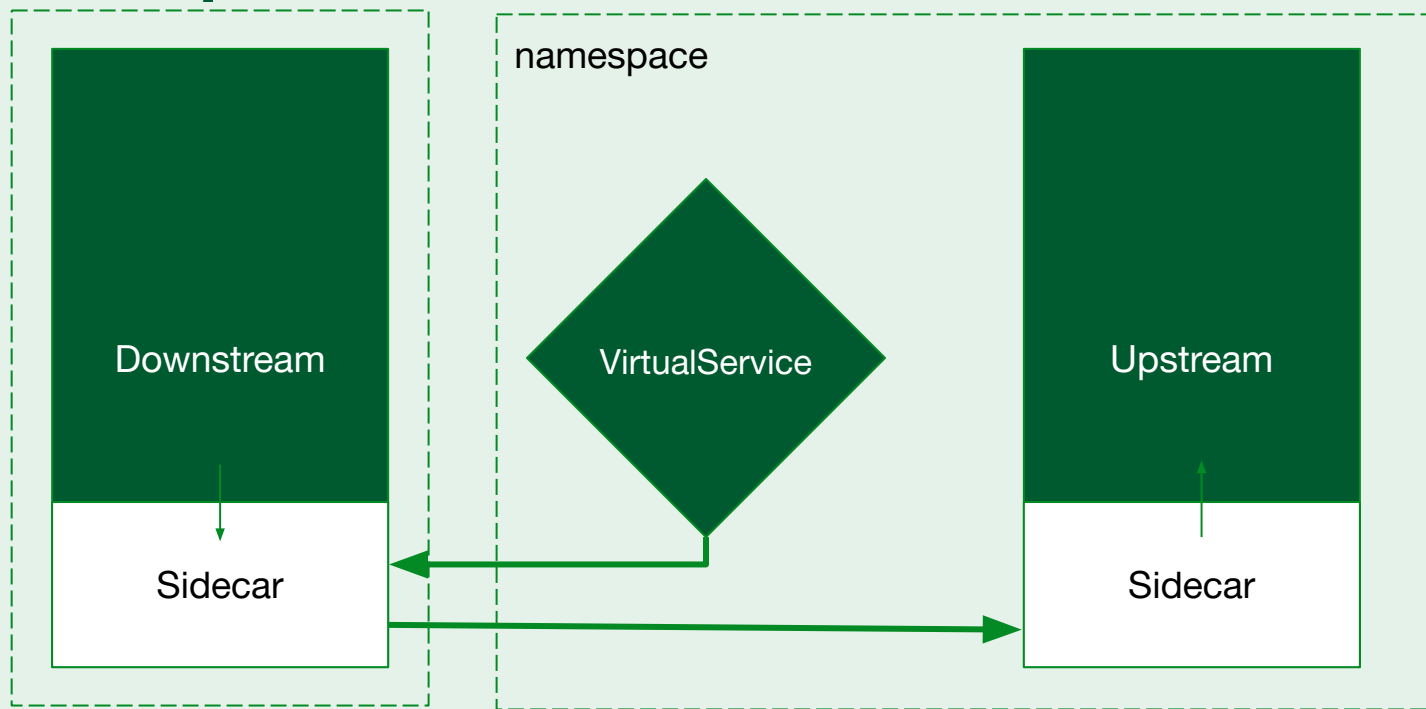
How you think about it



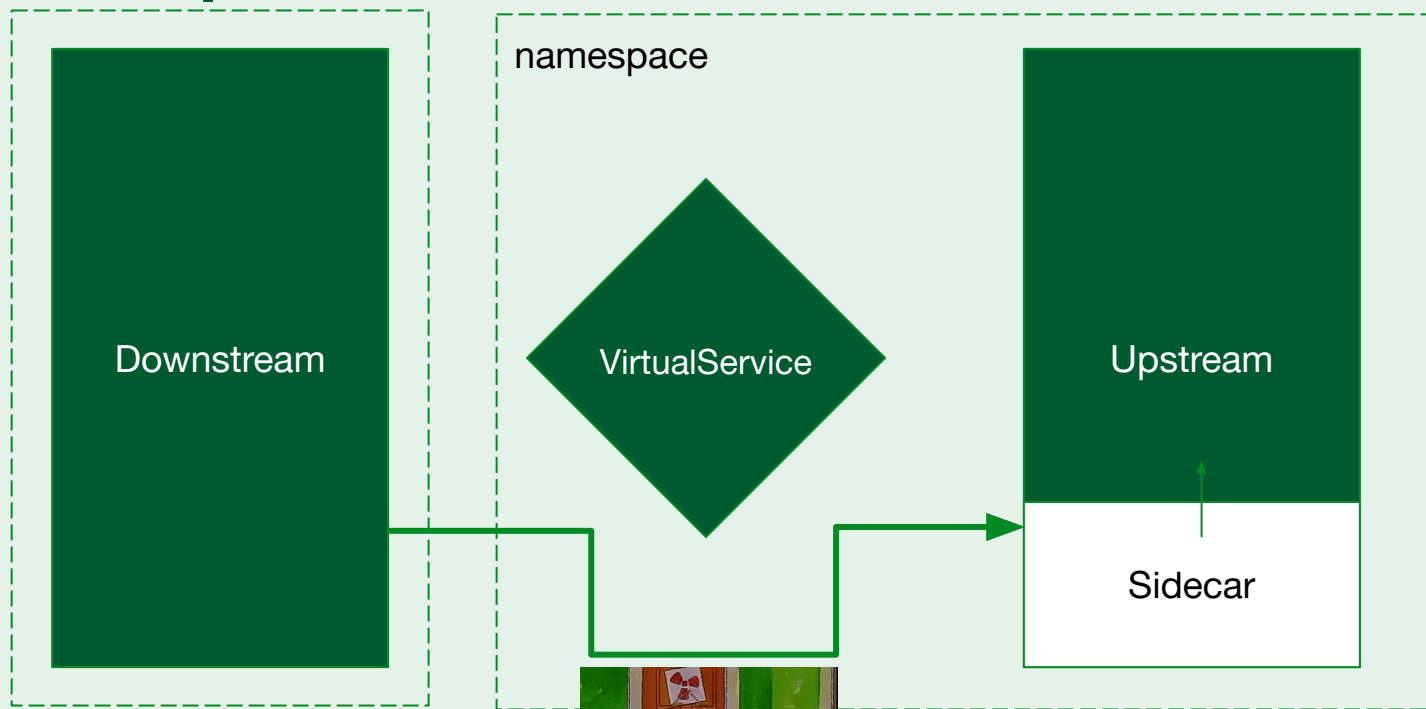
How you think about it



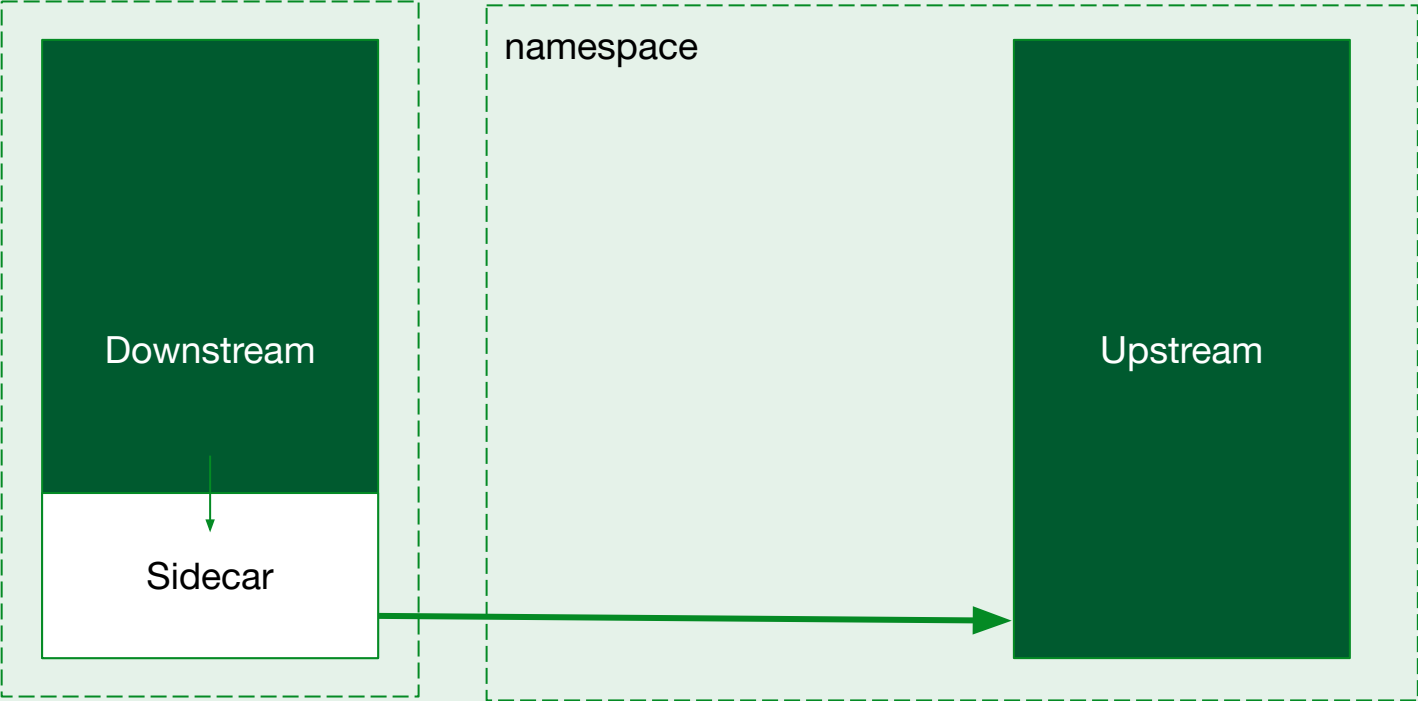
How It's Implemented



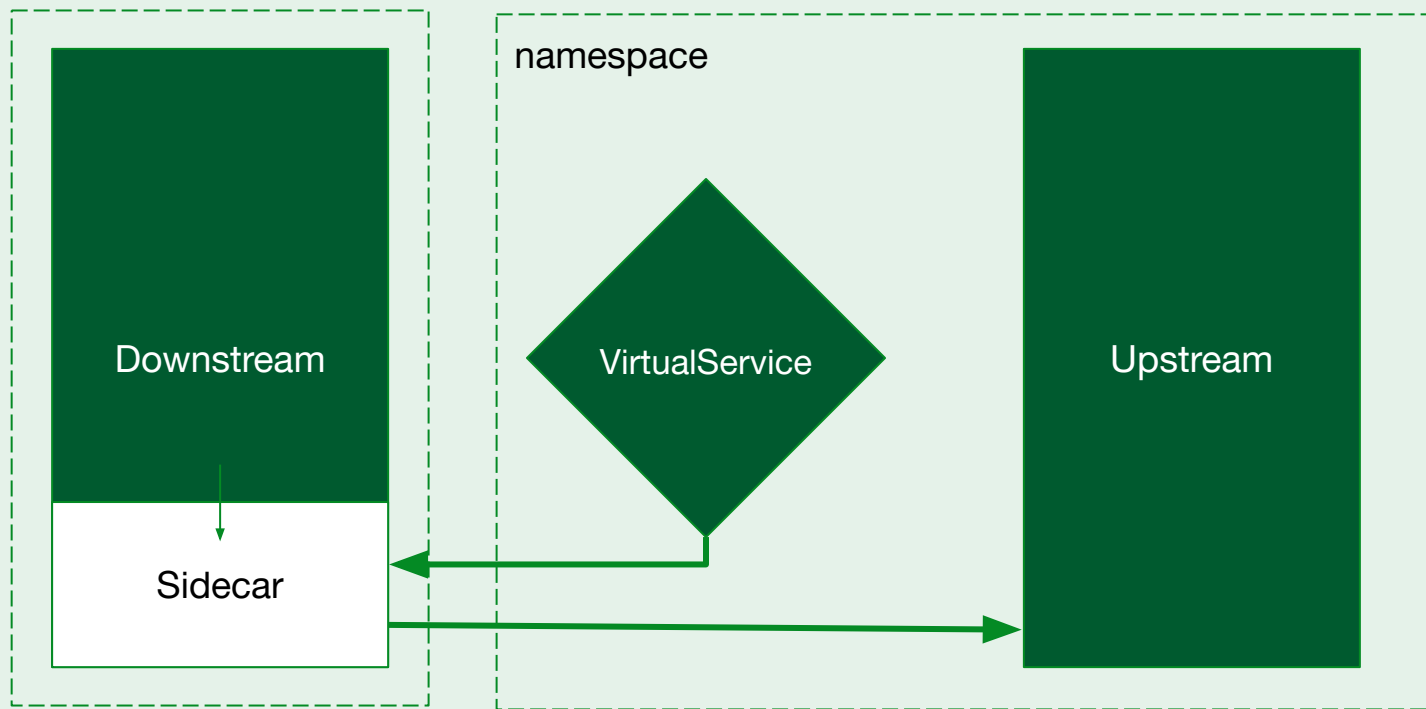
How It's Implemented



Downstream Only



Who knows!



It's All Happening Downstream

Retries

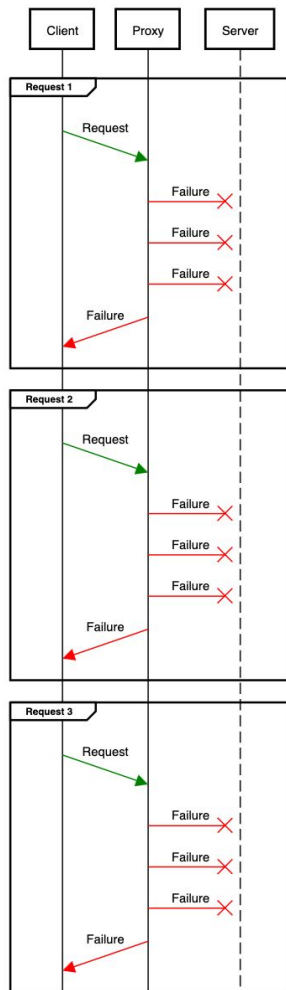
Before: Configured in Client (Downstream)

```
config :service_mesh_http_owl,  
  retry_count: 3,  
  retry_start: 250
```

Now: Configured in VirtualService (Upstream)

```
retries:  
  attempts: 3  
  retryOn: >=  
    5xx,  
    connect-failure,  
    retriabile-4xx,  
  perTryTimeout: 2s  
  backoff: 3s
```

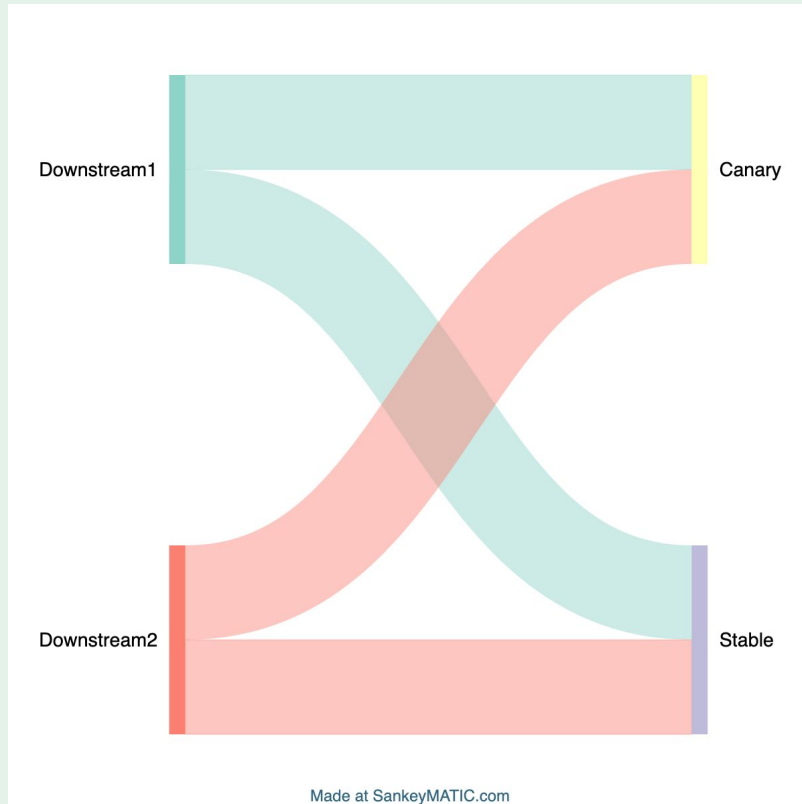
Don't Nest Retries!



Canaries



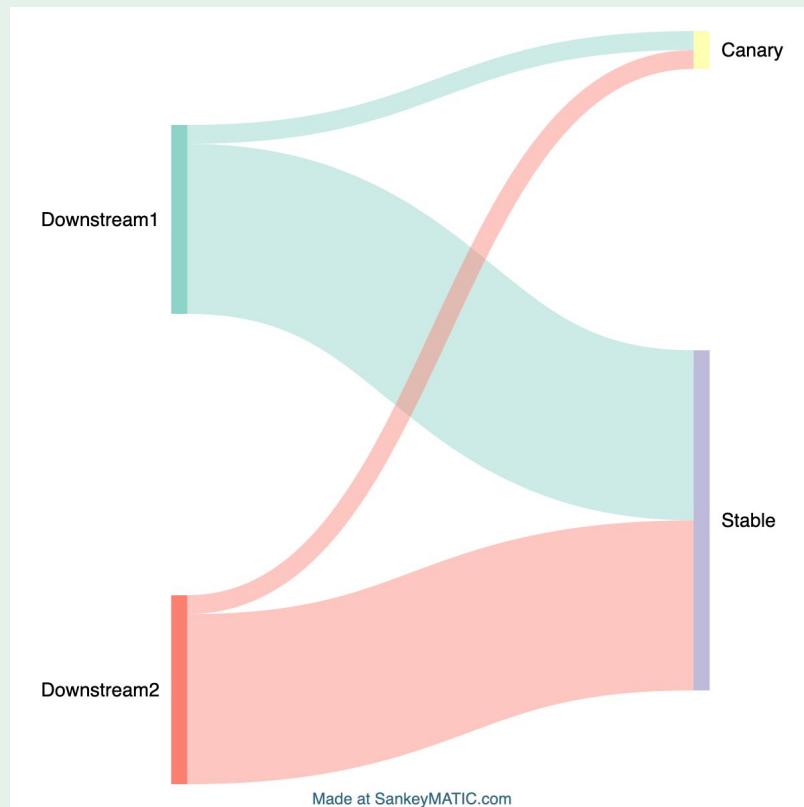
Canaries



50%

50%

Canaries



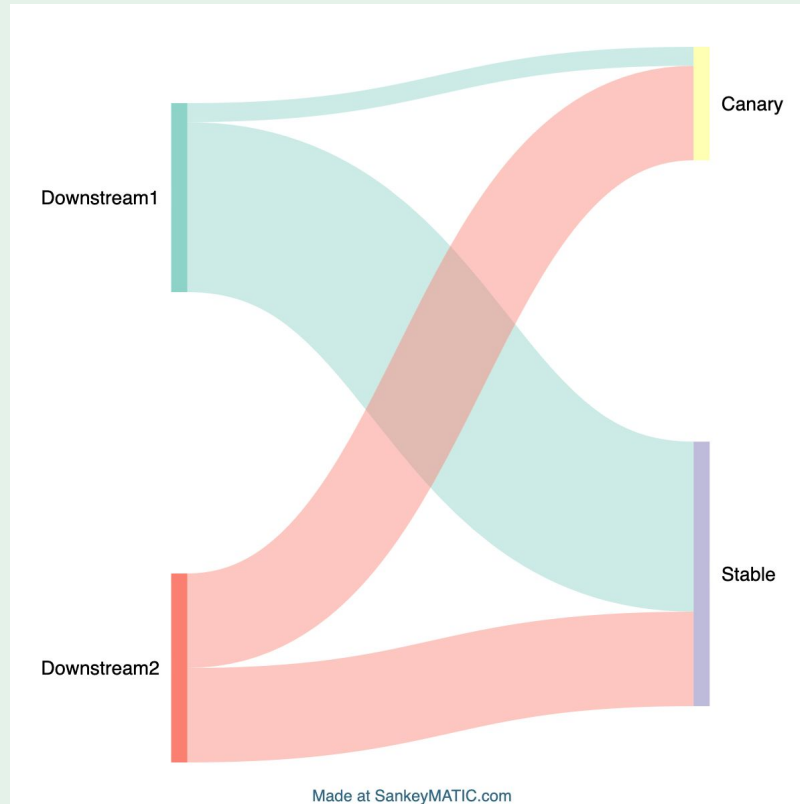
10%

90%

Canaries

mesh

no mesh



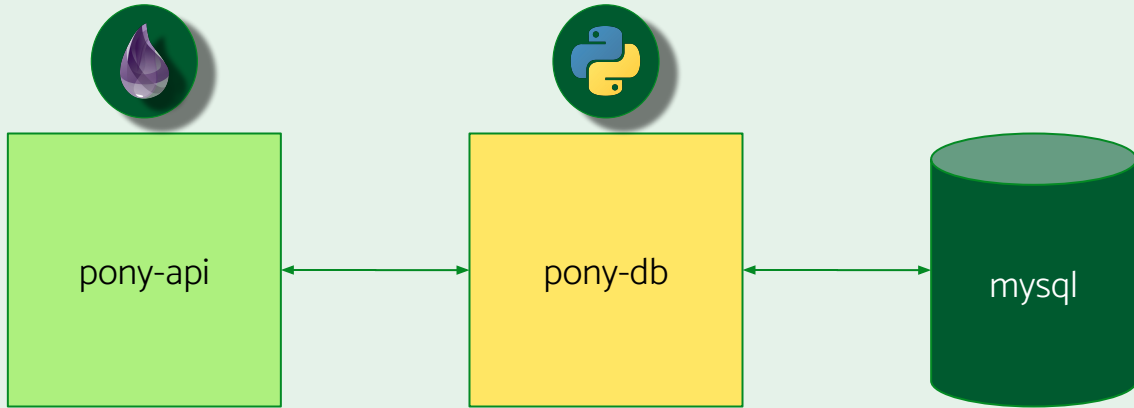
30%

70%

PagerDuty

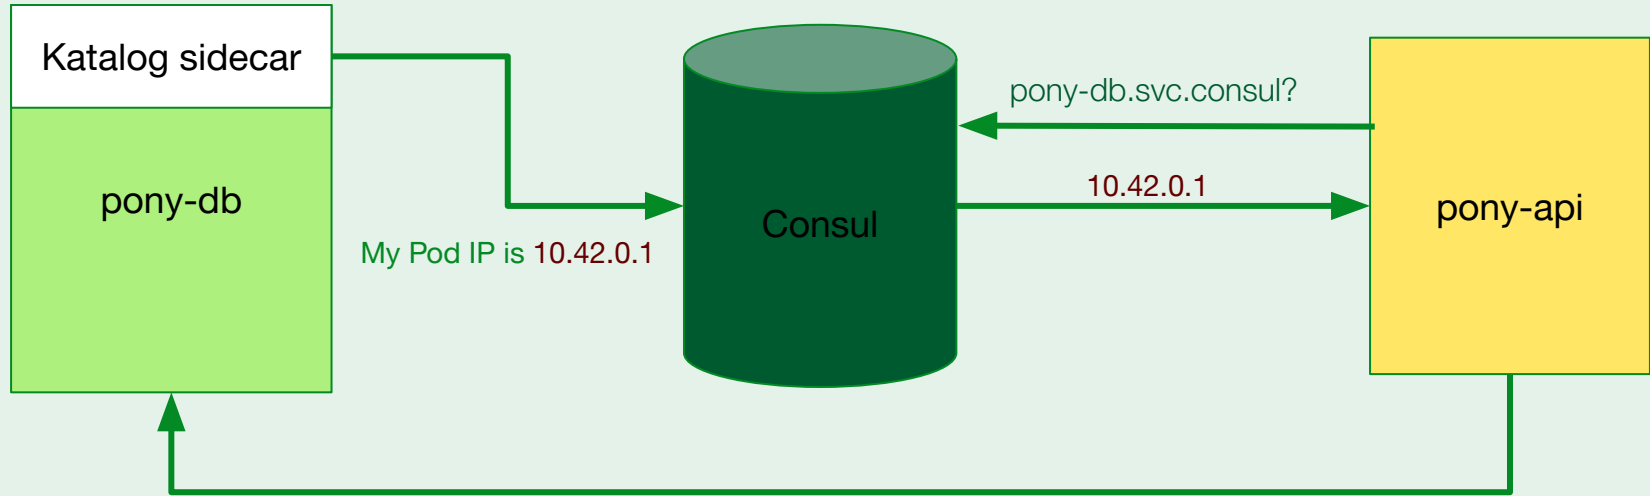
How we migrated

(how do you eat an elephant?)



PD's most critical service

Consul Discovery



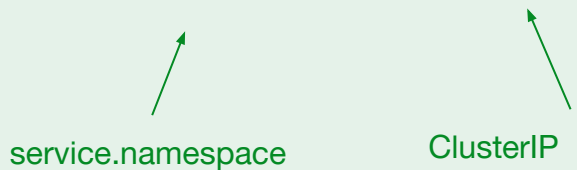
How does Istio route traffic?

VHOST NAME

pony-db.pony-db.svc.cluster.local:80

DOMAINS

pony-db.pony-db, 10.43.223.228



- ✓ `http://pony-db.pony-db`
- ✓ `http://pony-db.svc.cluster.local`
- ✓ `http://10.43.223.228`

Can we just match the domain?

```
apiVersion: networking.istio.io/v1
```

```
kind: VirtualService
```

```
metadata:
```

```
  name: pony-db
```

```
spec:
```

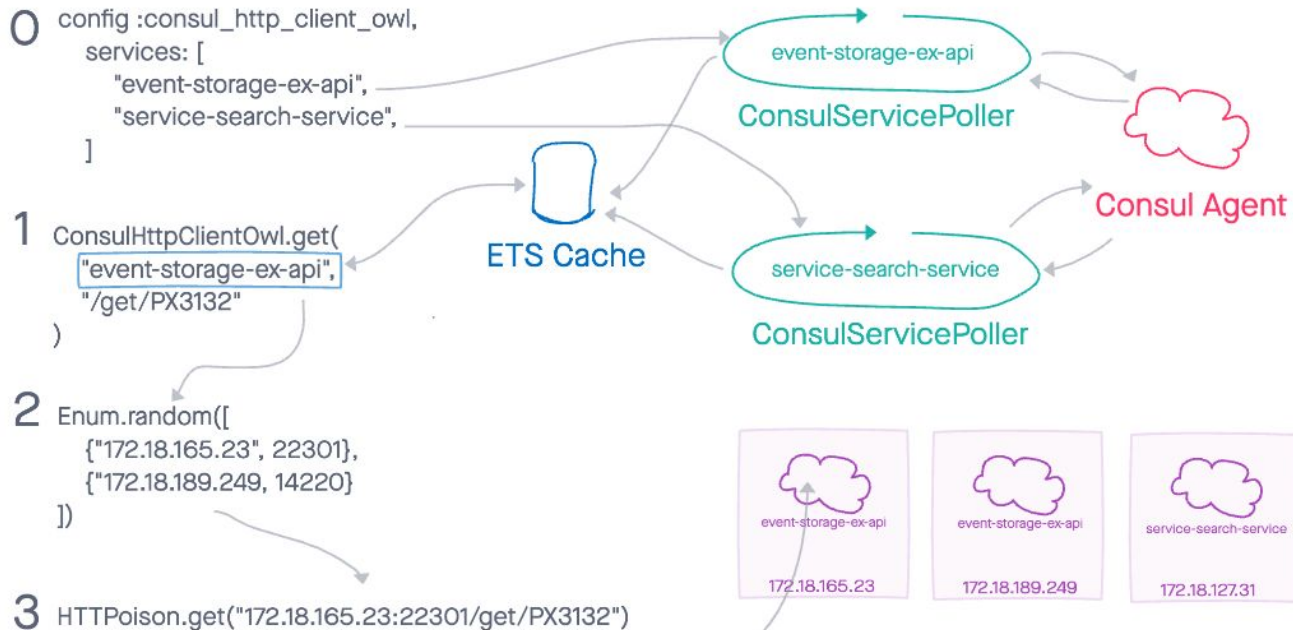
```
  hosts:
```

```
    - pony-db.service.consul
```

How does Istio route traffic?

- ✓ `http://pony-db.pony-db`
- ✓ `http://pony-db.svc.cluster.local`
- ✓ `http://10.43.223.228`
- ✓ `http://pony-db.service.consul`

Not just DNS



How does Istio route traffic?

- ✓ `http://pony-db.pony-db`
- ✓ `http://pony-db.svc.cluster.local`
- ✓ `http://10.43.223.228`
- ✓ `http://pony-db.service.consul`
- ✗ `http://10.42.0.1`

Two Phase Migration

Step 1: Kubedns

Two Phase Migration

Step 1: Kubedns

Also update dependencies, change wrapper libraries, add a new package repository, replace names...

Setup

Using a Script

Discovery Conversion

Dependencies

Updating mix.exs and mix.lock

Base HTTP Client

Client Wrappers

Adapting a new helper library

Configuration

Unit Test Configuration

Local Development Configuration

Library-Specific Configuration

permissions_client

entitlements_client

PdClient

sps-client

sss-client and rls-client

multi-region-http-client

Code Changes

Referenced Module

Services by Name and

Namespace

Tests

Tips

Sending the PR

Enabling the Sidecar

Joining the service mesh

Monitoring

AI?

sed: the original AI

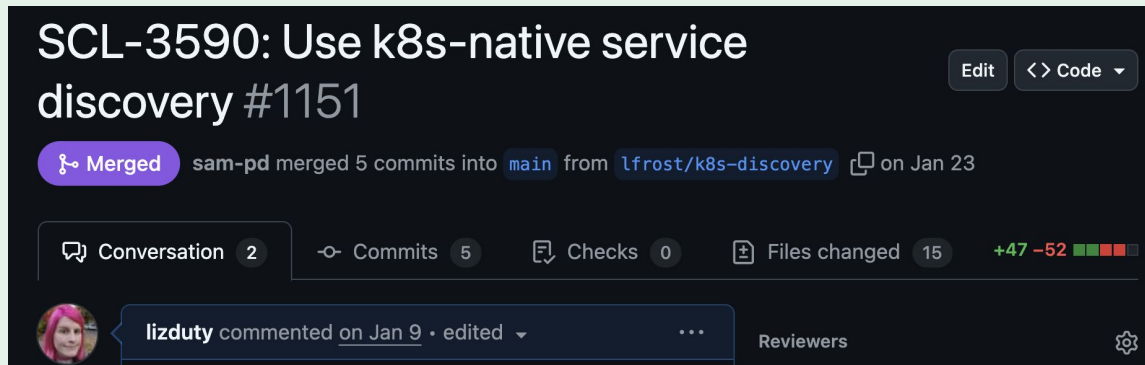
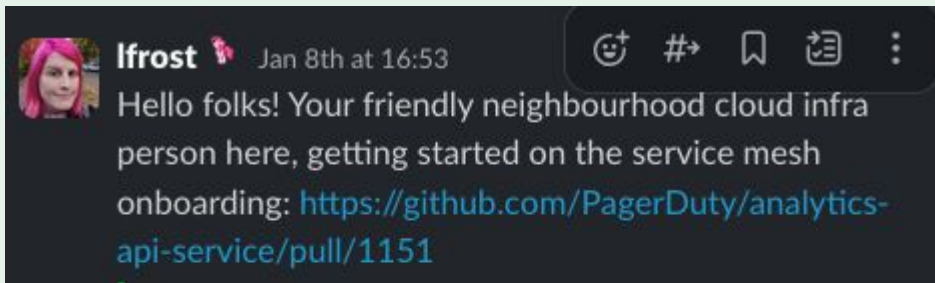
```
('ConsulHttpClientOwl', 'ServiceMeshHttpOwl'),
('ConsulServiceLookupOwl', 'ServiceMeshHttpOwl'),
('consul_http_client_owl_mod', 'http_client_mod'),
('consul_http_client_owl', 'service_mesh_http_owl'),
('consul_service_lookup_owl', 'service_mesh_http_owl'),
('consul_name', 'service_name'),
('PdClientMock', 'PdWebClientMock'),
('PdClient.Mock', 'PdWebClientMock'),
('PdClient.Behaviour', 'ServiceMeshHttpOwl.PdWebClientBehaviour'),
('PdClient', 'ServiceMeshHttpOwl.PdWebClient')
```

Away Team Model

aka “white glove” or “we code so you don’t have to”

Away Team Model

- Divided up by language
- Concentrates experience
- Many hands make light(er) work
- Preserves social capital



Two Phase Migration

```
+ mesh:  
+   enabled: true
```

The screenshot shows a GitHub pull request interface. At the top, the title is "SCL-3676: Enable the service mesh #237" with "Edit" and "Code" buttons. Below the title, a purple "Merged" badge is followed by the text "LuckDragon82 merged 1 commit into main from lfrost/enable-mesh on Feb 5". A progress bar shows "Conversation 15", "Commits 1", "Checks 0", and "Files changed 2" with a net change of "+9 -0".

A comment from user "lizduty" is visible, dated "Feb 5". The comment text reads: "We've enabled core-dns discovery, so now it's time for actual mesh onboarding. Please look over this carefully, it is not a PR from your team, there may be mistakes!". Below the comment, a section titled "Your to do list:" is partially visible.

On the right side of the comment, there are sections for "Reviewers" (with a green checkmark), "Assignees" (with the text "No one—assign yourself"), and "Labels" (with the text "None yet").

Helm Chart

```
mesh:  
  enabled: false  
  access_logs:  
    enabled: false  
  envoy_dns_cache:  
    disable: false  
  settings:  
    max_request_headers_kb: 120
```

Access Logs

```
spec:  
  accessLogging:  
    - disabled: false  
      providers:  
        - name: full  
  selector:  
    matchLabels:  
      app.kubernetes.io/instance: pony-api  
      app.kubernetes.io/name: pony-api  
      app.pagerduty.com/env: production  
      app.pagerduty.com/service-region: prod  
      app.pagerduty.com/short-env: prod
```

max_header_size?

```
spec:  
  configPatches:  
    - applyTo: NETWORK_FILTER  
      match:  
        context: SIDECAR_INBOUND  
        listener:  
          filterChain:  
            filter:  
              name: envoy.filters.network.http_connection_manager  
      patch:  
        operation: MERGE  
        value:  
          typed_config:  
            '@type': >-
```

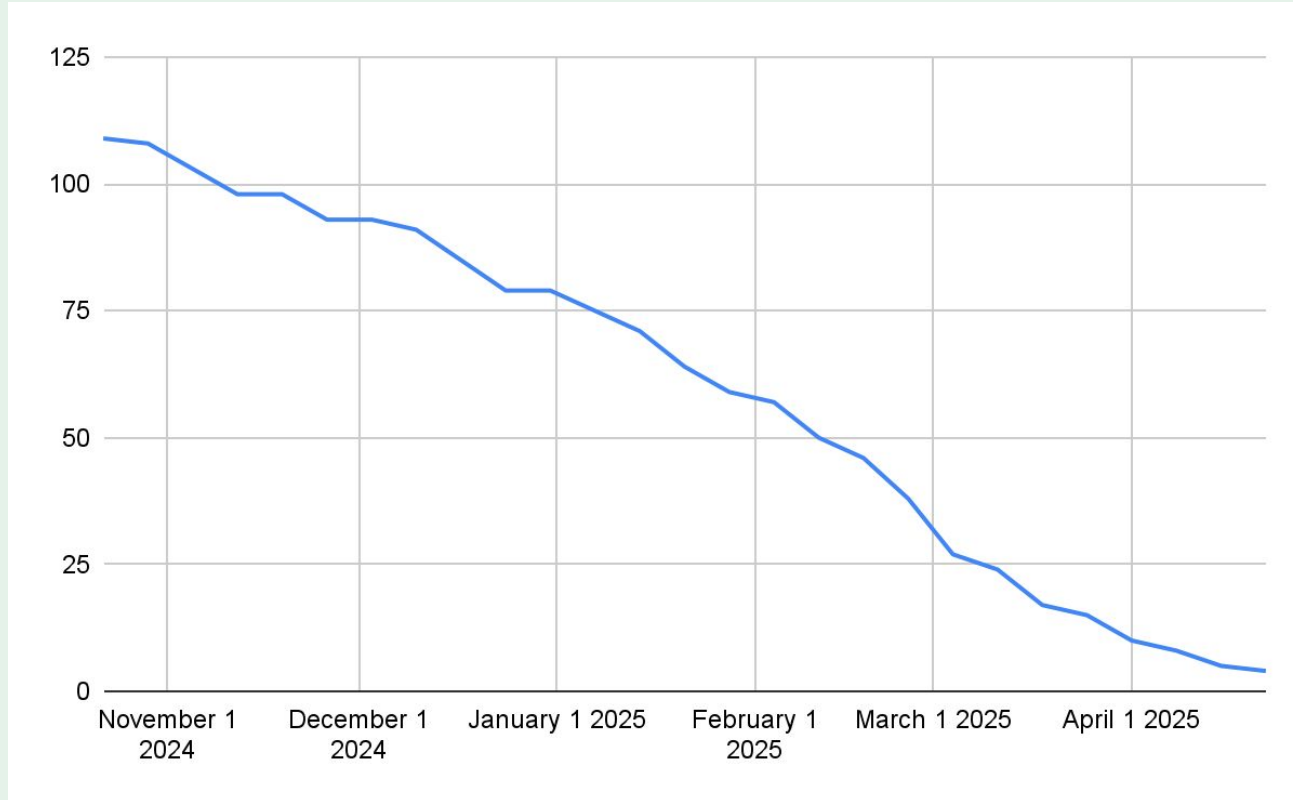
[type.googleapis.com/envoy.extensions.filters.network.http_connection_manager.v3.HttpConnectionMa
nager](https://type.googleapis.com/envoy.extensions.filters.network.http_connection_manager.v3.HttpConnectionManager)

```
      max_request_headers_kb: 128  
  workloadSelector:  
    labels:  
      app.kubernetes.io/name: advance-configuration-service
```

Tracking Progress

	<i>Planned Migration Timeline</i>	<i>App type</i>	COUNT of NAM		
8					
9					
0	+ 13 - January Total		31		
1	+ 08 - August Total		10		
2	+ 09 - September Total		5		
3	+ 10 - October Total		51		
4	+ 11 - November Total		38		
5	+ 12 - December Total		30		
6	Grand Total		165		

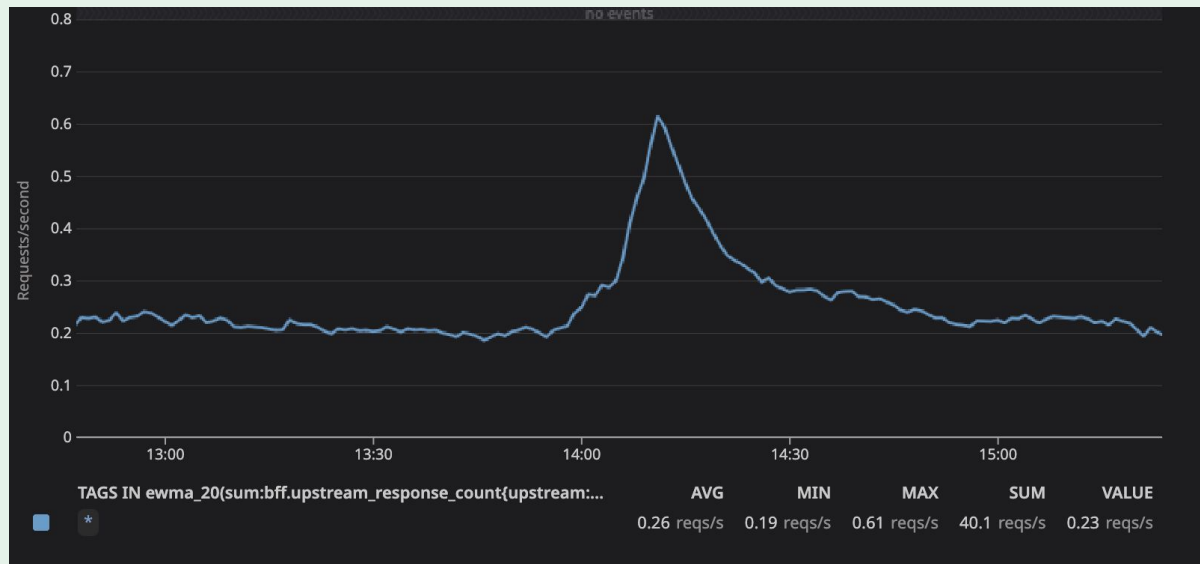
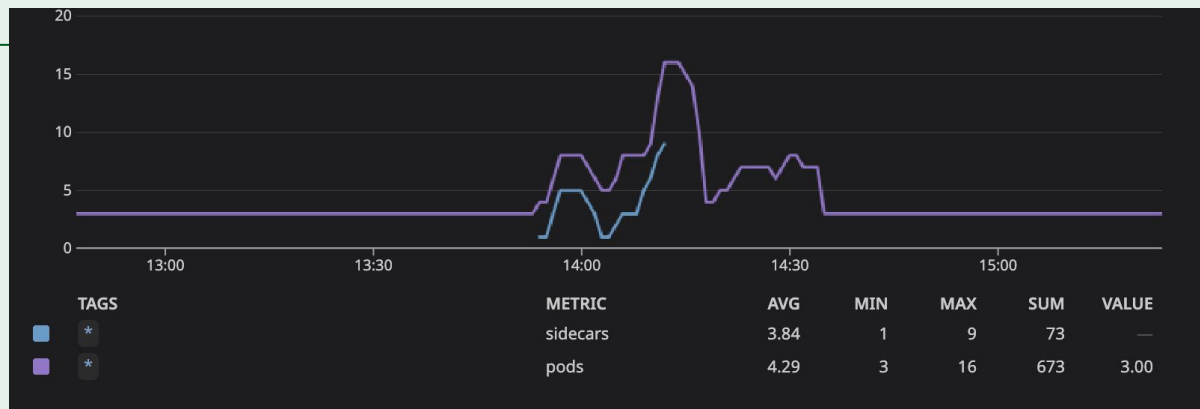
slow and steady



The Big Rock



The Symptom



They're all timeouts

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:22 UTC

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:18 UTC

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:17 UTC

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:17 UTC

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:16 UTC

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:15 UTC

Rack::Timeout::RequestTimeoutException

Request ran for longer than 44000ms

4 months ago · May 8, 21:12:15 UTC

What didn't work

- Both client and server on mesh
- Isolated Kubernetes node
- More memory/CPU
- conntrack allowances
- Any sort of reproduction in staging

YOU!



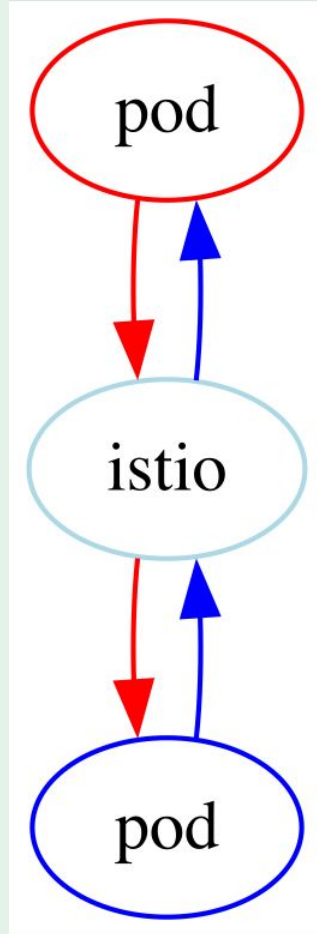
Ongoing work!

- TProxy?
- Nftables / conntrack
- Combinatorics

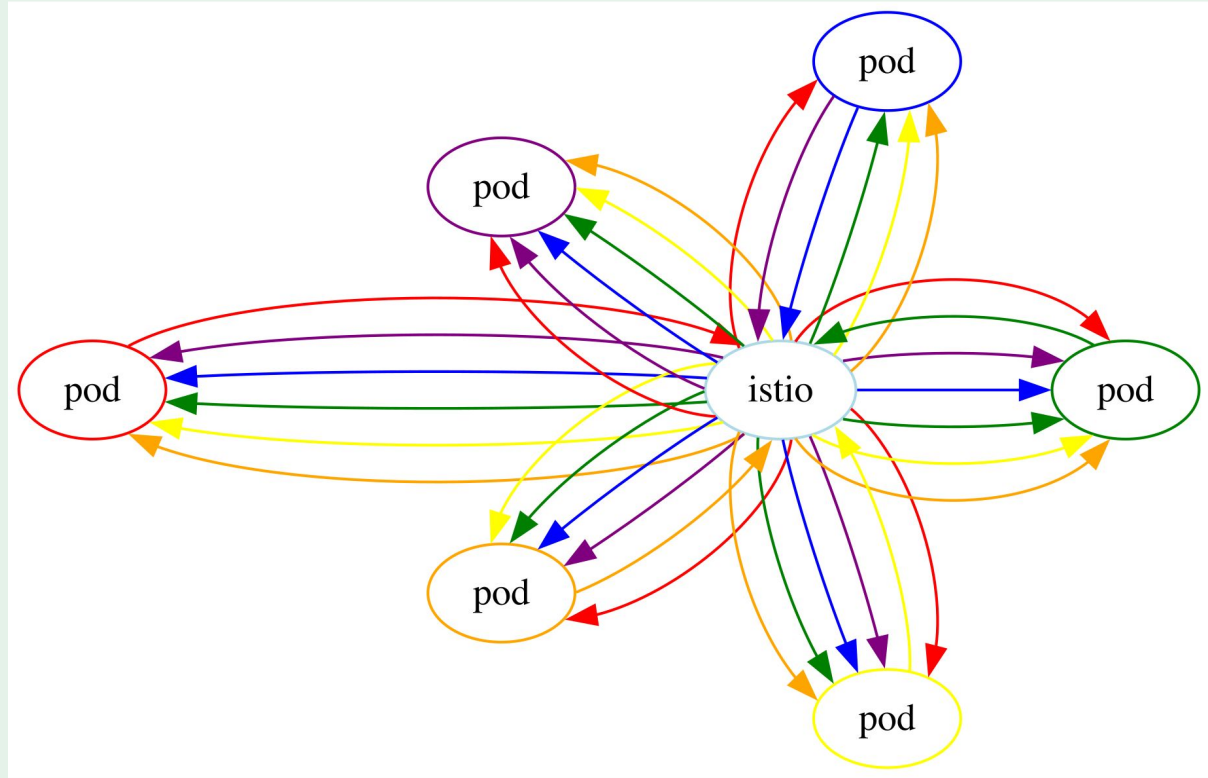
Ongoing work!

- TProxy?
- Nftables / conntrack
- **Combinatorics**

Two pods

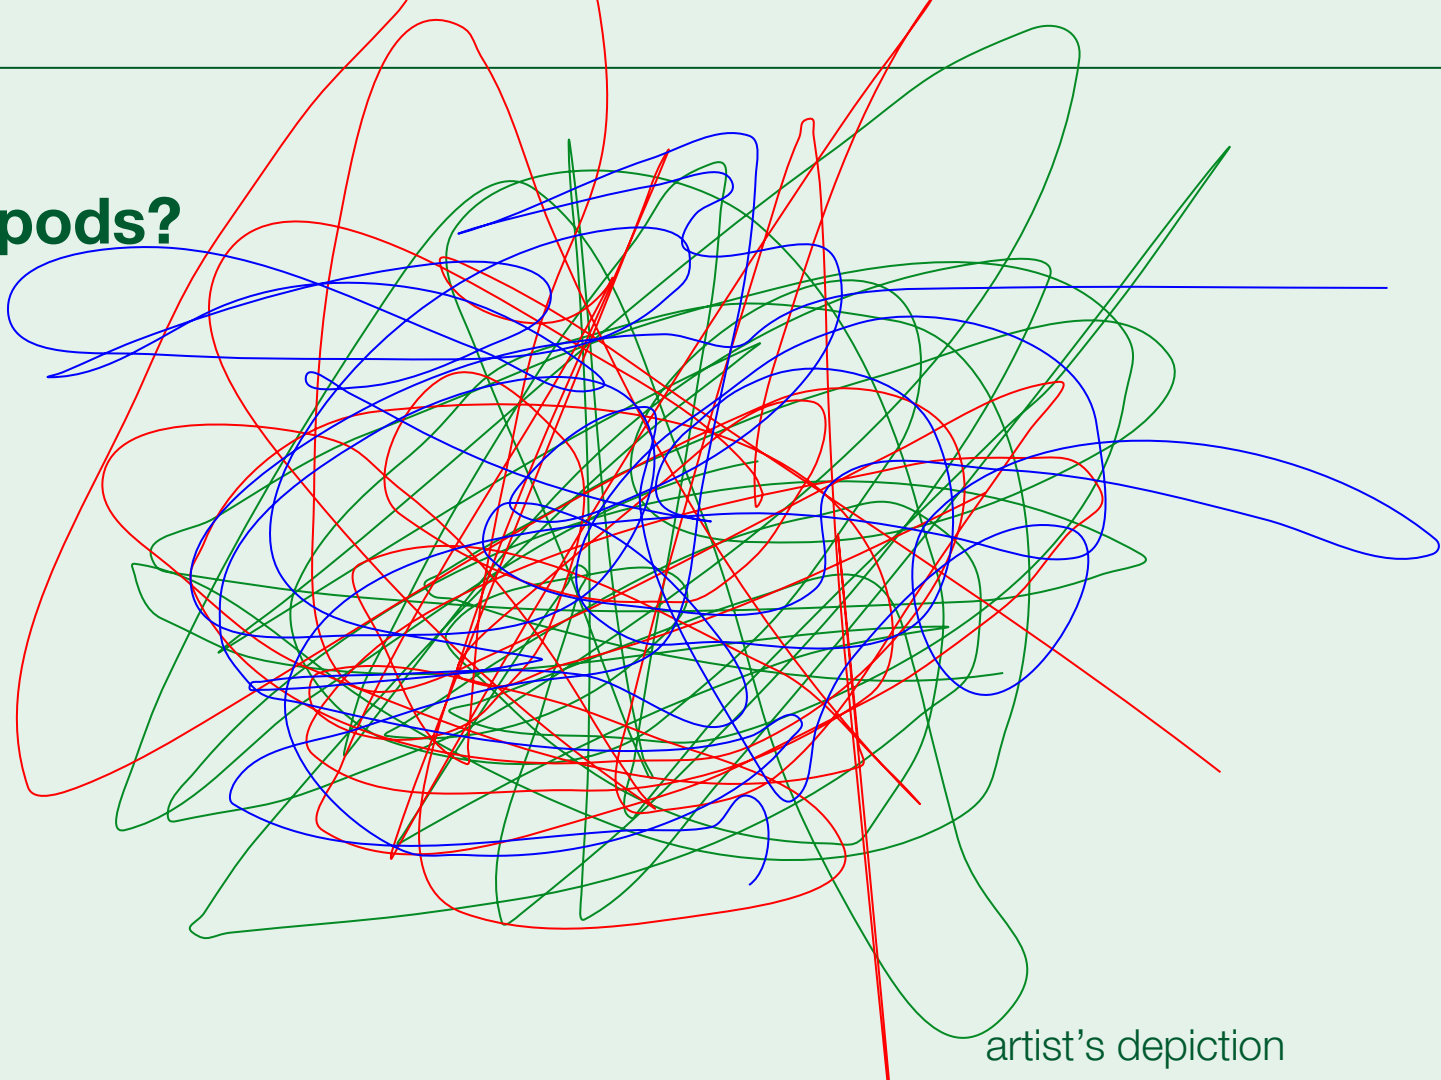


Six pods



4000+ pods?

4000+ pods?



artist's depiction

PagerDuty

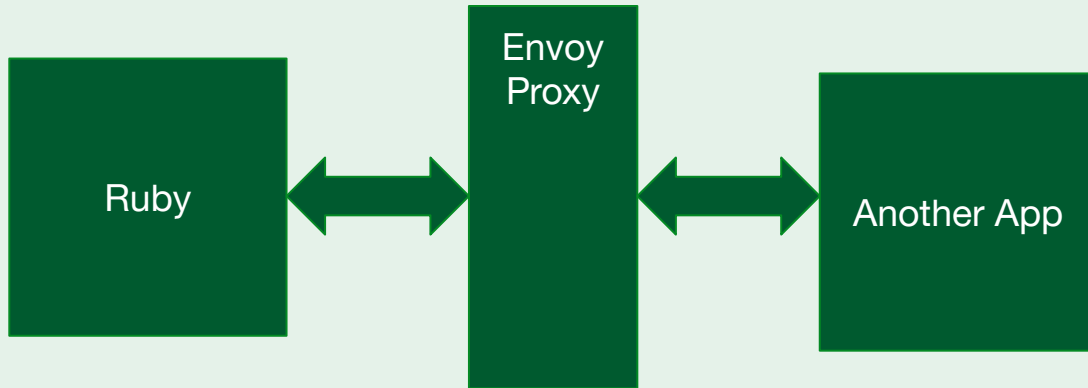
Random Gotchas

legos on the floor

Not everything is http

```
apiVersion: security.istio.io/v1beta1
kind: PeerAuthentication
metadata:
  name: disable-mtls-for-smtp
  namespace: postfix-email-ingestion
spec:
  selector:
    matchLabels:
      app.kubernetes.io/name:
postfix-email-ingestion
  mtls:
    mode: UNSET
  portLevelMtls:
    25:
      mode: DISABLE
```

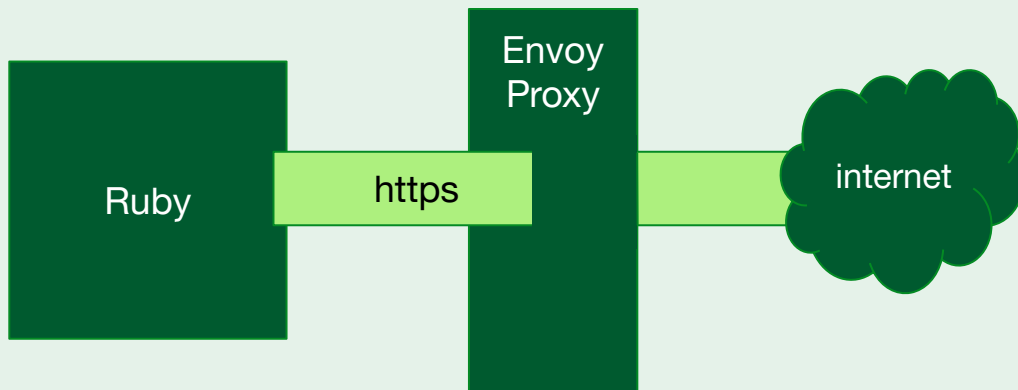
Normal Envoy flow



Proxied TLS issues

```
SSL_write() error: error:1409E10F:SSL routines:ssl3_write_bytes:bad length
```

Typhoeus (libcurl) to Faraday(net/http)



No native rate limiting :(

```
apiVersion: networking.istio.io/v1alpha3
kind: EnvoyFilter
metadata:
  name: filter-ratelimit
  namespace: istio-system
spec:
  workloadSelector:
    # select by label in the same namespace
    labels:
      istio: ingressgateway
  configPatches:
    # The Envoy config you want to modify
    - applyTo: HTTP_FILTER
      match:
        context: GATEWAY
        listener:
          filterChain:
            filter:
              name: "envoy.filters.network.http_connection_manager"
```

Max header Size



431

Request Header Fields Too Large

Envoy Filters

```
spec:
  configPatches:
    - applyTo: NETWORK_FILTER
      match:
        context: SIDECAR_INBOUND
        listener:
          filterChain:
            filter:
              name: envoy.filters.network.http_connection_manager
      patch:
        operation: MERGE
        value:
          typed_config:
            '@type': >-
```

```
type.googleapis.com/envoy.extensions.filters.network.http_connection_manager.
v3.HttpConnectionManager
```

```
    max_request_headers_kb: 128
```

```
workloadSelector:
```

```
  labels:
```

```
    app.kubernetes.io/name: advance-configuration-service
```

Where do they go?

- Just read istio github issues until someone solves it

```
$ istioctl proxy-config clusters $POD -o json | wc -l
```

```
4655
```

jq?

```
.configs[] |  
select( ."@type" |  
contains("Listeners") ) |  
.dynamic_listeners[] |  
select (.name ==  
"virtualInbound").active_state.listener.filter_chains[].filters[] |  
select (.name == "envoy.filters.network.http_connection_manager")
```

what could possibly go wrong –

```
panic: runtime error: slice bounds out of range [-2:]
```

```
goroutine 19414 [running]:
```

```
github.com/envoyproxy/go-control-plane/envoy/config/core/v3.(*SocketAddress).MarshalToSizedBufferVTStrict(0xc0070b7730, {0xc0013f0000, 0x8, 0x15be})
```

```
github.com/envoyproxy/go-control-plane@v0.12.1-0.20240415211714-57c85e1829e6/envoy/config/core/v3/address_vtproto.pb.go:191 +0x415
```

```
github.com/envoyproxy/go-control-plane/envoy/config/core/v3.(*Address_SocketAddress).MarshalToSizedBufferVTStrict(0xa47?, {0xc0013f0000, 0x8, 0x0?})
```

```
github.com/envoyproxy/go-control-plane@v0.12.1-0.20240415211714-57c85e1829e6/envoy/config/core/v3/address_vtproto.pb.go:507 +0x94
```

```
github.com/envoyproxy/go-control-plane/envoy/config/core/v3.(*Address).MarshalToSizedBufferVTStrict(0xc00608cc80, {0xc0013f0000, 0x8, 0x15be})
```

```
github.com/envoyproxy/go-control-plane@v0.12.1-0.20240415211714-57c85e1829e6/envoy/config/core/v3/address_vtproto.pb.go:490 +0x19a
```

```
github.com/envoyproxy/go-control-plane/envoy/config/listener/v3.(*Listener).MarshalToSizedBufferVTStrict(0xc00304ad00, {0xc0013f0000, 0x15be, 0x15be})
```

```
github.com/envoyproxy/go-control-plane@v0.12.1-0.20240415211714-57c85e1829e6/envoy/config/listener/v3/listener_vtproto.pb.go:782 +0x12e6
```

```
github.com/envoyproxy/go-control-plane/envoy/config/listener/v3.(*Listener).MarshalVTStrict(0xc00304ad00)
```

Resource Utilisation

Can't autoscale sidecars :(

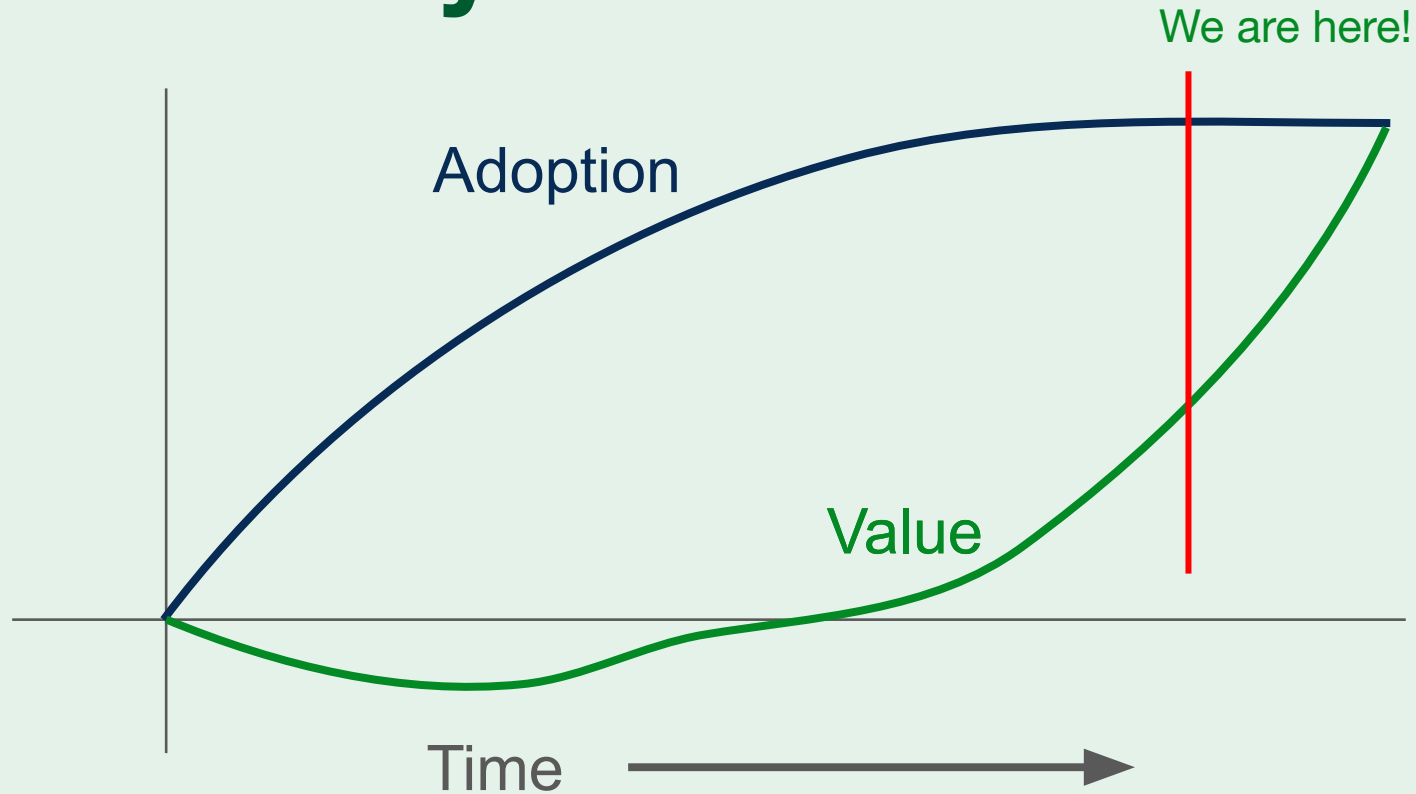
What is the problem?

The service mesh sidecar (Envoy proxy) is exhibiting high resource usage, which can lead to instability. You may be here because the service mesh sidecar monitoring module [service_mesh_of_terraform-pd-service](#) has referred you to this runbook.

CronJob support required a Kubernetes Upgrade

- Native Sidecar supported added in k8s 1.29
- Stable now!
- Default in Istio!
- Don't worry about it!!

The Reality



Thank you

Liz Frost
lfrost@pagerduty.com

Booth 26
Shifts: Weds / Thurs morning