

Comprehensive Container-Based Service Monitoring with Kubernetes and Istio

SREcon Asia Australia 2018-06-06
Fred Moyer

Monitoring Nerd

[@phredmoyer](#)

Developer Evangelist

[@Circonus](#) / [@IRONdb](#)

[@IstioMesh](#) Geek

Observability and
Statistics Dork



Talk Agenda


- ❑ Istio Overview
- ❑ Service Level Objectives
- ❑ RED Dashboard
- ❑ Histogram Telemetry
- ❑ Istio Metrics Adapter
- ❑ Asking the Right Questions

Istio.io

“An open platform to connect, manage, and secure microservices”




Happy Birthday!



K8S + Istio

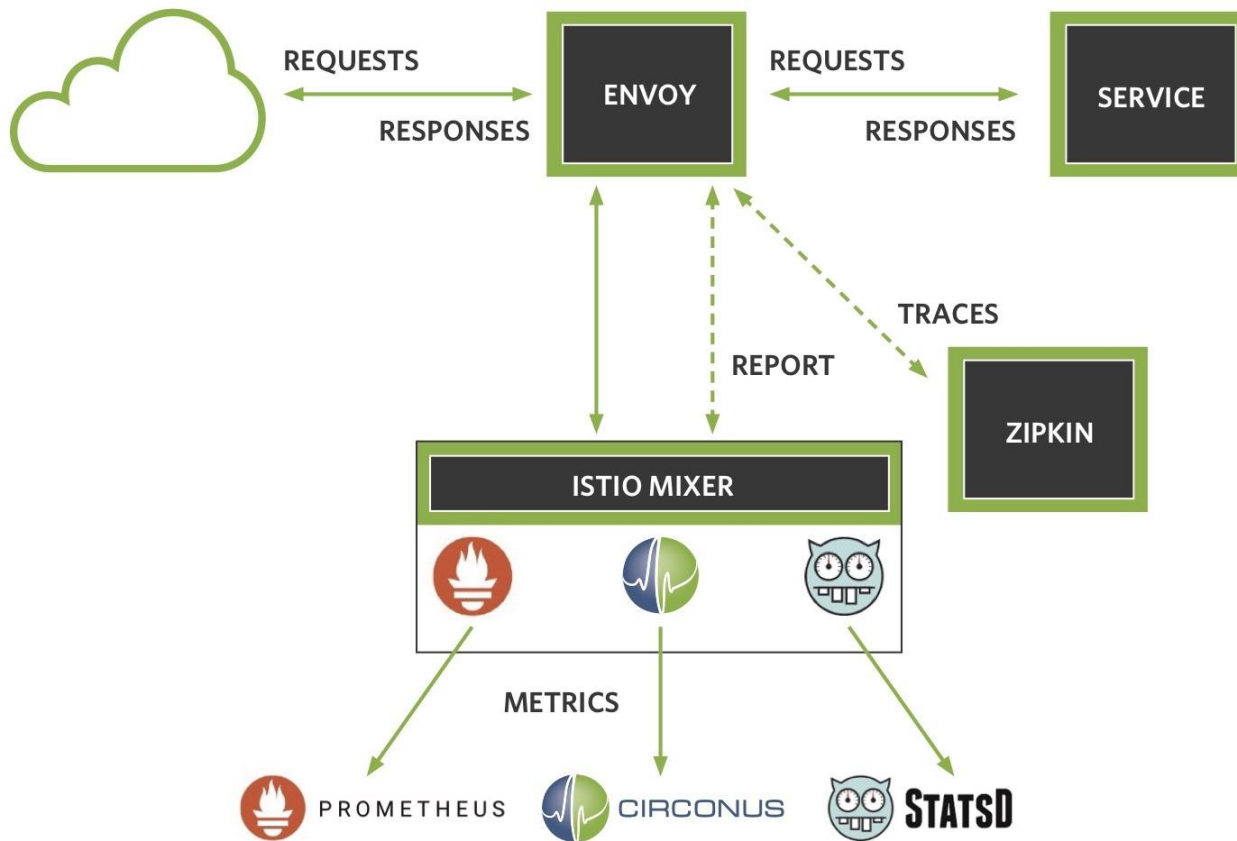


- Orchestration
- Deployment
- Scaling
- Data Plane



- Policy Enforcement
- Traffic Management
- Telemetry
- Control Plane

Istio Architecture



Istio GCP Deployment

Deployment name

installIstioRelease

[^ Less](#)

Baseline GKE Cluster config

GKE Cluster Name [?](#)

Zone [?](#)

Number of GKE nodes to run on [?](#)

Node Machine Type [?](#)

7.5 GB memory

[Customize](#)

Monitoring, Logging and Tracing

Metrics, Logs, and Traces [?](#)

Add optional Metrics, Logs, and Traces related plugins to the cluster

- Enable Prometheus for metrics/logs collection [?](#)
- Enable Grafana for metrics display [?](#)
- Enable Zipkin for tracing [?](#)
- Enable ServiceGraph for deployment visualization [?](#)

Security

Security [?](#)

Add optional Security related plugins to the cluster

- Enable Automatic Istio sidecar injection [?](#)
- Enable mutualTLS authentication [?](#)

Install Applications

- Add BookInfo Sample Application [?](#)

[Deploy](#)

Istio Sample App

```
$ istioctl create -f apps/bookinfo.yaml
```

Istio Sample App

BookInfo Sample Sign in

The Comedy of Errors

Wikipedia Summary: The Comedy of Errors is one of **William Shakespeare's** early plays. It is his shortest and one of his most farcical comedies, with a major part of the humour coming from slapstick and mistaken identity, in addition to puns and word play.

Book Details

Paperback:
200 pages

Publisher:
PublisherA

Language:
English

ISBN-10:
1234567890

ISBN-13:
123-1234567980

An extremely entertaining play by Shakespeare. The slapstick humour is refreshing!

— Reviewer1 *Affiliation1*

★★★★★

Absolutely fun and entertaining. The play lacks thematic depth when compared to other plays by Shakespeare.

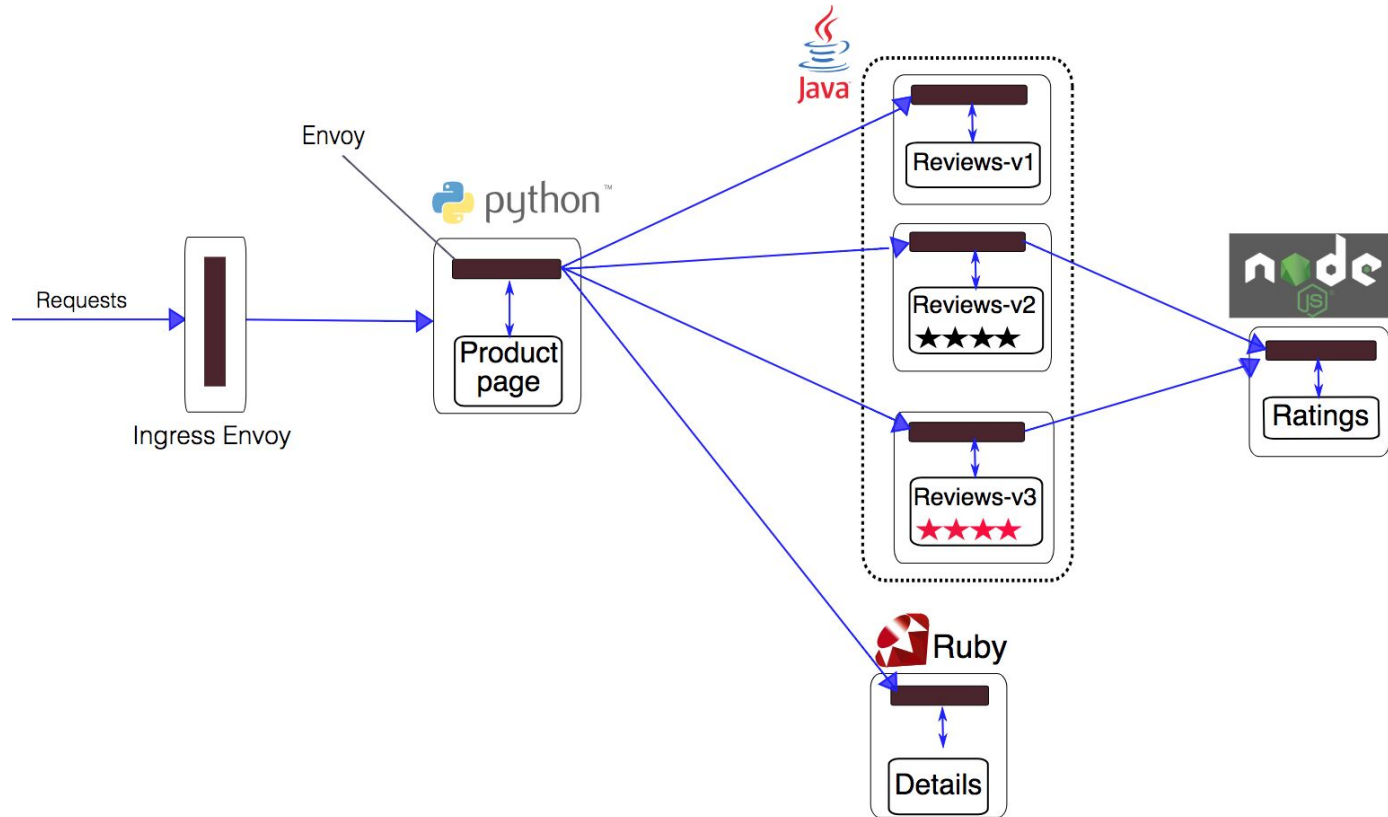
— Reviewer2 *Affiliation2*

★★★★☆

```
precedence: 1
route:
- tags:
  version: v1
  weight: 100
Go refresh the page
```



Istio Sample App



Istio Sample App

```
kind: Deployment
metadata:
  name: ratings-v1
spec:
  replicas: 1
  template:
    metadata:
      labels:
        app: ratings
        version: v1
    spec:
      containers:
      - name: ratings
        image: istio/examples-bookinfo-ratings-v1
        imagePullPolicy: IfNotPresent
        ports:
        - containerPort: 9080
```

Istio Sample App

```
$ istioctl create -f apps/bookinfo/route-rule-reviews-v2-v3.yaml
```

```
type: route-rule
```

```
spec:
```

```
  name: reviews-default
```

```
  destination: reviews.default.svc.cluster.local
```

```
  precedence: 1
```

```
  route:
```

```
    - tags:
```

```
      version: v2
```

```
      weight: 80
```

```
    - tags:
```

```
      version: v3
```

```
      weight: 20
```

Istio K8s Services

```
> kubectl get services
```

NAME	CLUSTER-IP	EXTERNAL-IP	PORT (S)	AGE
details	10.0.0.31	<none>	9080/TCP	6m
kubernetes	10.0.0.1	<none>	443/TCP	7d
productpage	10.0.0.120	<none>	9080/TCP	6m
ratings	10.0.0.15	<none>	9080/TCP	6m
reviews	10.0.0.170	<none>	9080/TCP	6m

Istio K8s App Pods

```
> kubectl get pods
```

NAME	READY	STATUS	RESTARTS	AGE
details-v1-1520924117	2/2	Running	0	6m
productpage-v1-560495357	2/2	Running	0	6m
ratings-v1-734492171	2/2	Running	0	6m
reviews-v1-874083890	2/2	Running	0	6m
reviews-v2-1343845940	2/2	Running	0	6m
reviews-v3-1813607990	2/2	Running	0	6m

Istio K8s System Pods

```
> kubectl get pods -n istio-system
```

NAME	READY	STATUS	RESTARTS	AGE
istio-ca-797dfb66c5	1/1	Running	0	2m
istio-ingress-84f75844c4	1/1	Running	0	2m
istio-egress-29a16321d3	1/1	Running	0	2m
istio-mixer-9bf85fc68	3/3	Running	0	2m
istio-pilot-575679c565	2/2	Running	0	2m
grafana-182346ba12	2/2	Running	0	2m
prometheus-837521fe34	2/2	Running	0	2m




Talk Agenda



- ✓ Istio Overview
- ❑ Service Level Objectives
- ❑ RED Dashboard
- ❑ Histogram Telemetry
- ❑ Istio Metrics Adapter
- ❑ Asking the Right Questions

Service Level Objectives

- **SLI** - Service Level Indicator
- **SLO** - Service Level Objective
- **SLA** - Service Level Agreement



Service Level Objectives



“SLIs drive SLOs which inform SLAs”

SLI - Service Level Indicator, a measure of the service that can be quantified

“95th percentile latency of homepage requests over past 5 minutes < 300ms”

Excerpted from
“SLIs, SLOs, SLAs, oh my!”
[@sethvargo](#) [@lizthegrey](#)

<https://youtu.be/tEylFyxbDLE>

“SLIs drive SLOs which inform SLAs”

SLO - Service Level Objective, a target
for Service Level Indicators

*“95th percentile homepage SLI will
succeed 99.9% over trailing year”*

Excerpted from
“SLIs, SLOs, SLAs, oh my!”
[@sethvargo](#) [@lizthegrey](#)

<https://youtu.be/tEylFyxbDLE>

“SLIs drive SLOs which inform SLAs”

SLA - Service Level Agreement, a legal agreement between a customer and a service provider based on SLOs

“Service credits if 95th percentile homepage SLI succeeds less than 99.5% over trailing year”

Excerpted from
“SLIs, SLOs, SLAs, oh my!”
[@sethvargo](#) [@lizthegrey](#)

<https://youtu.be/tEylFyxbDLE>

Talk Agenda

- ✓ Istio Overview
- ✓ Service Level Objectives
- ❑ RED Dashboard
- ❑ Histogram Telemetry Collection
- ❑ Istio Metrics Adapter
- ❑ Asking the Right Questions

Emerging Standards

- **USE**
 - Utilization, Saturation, Errors
 - Introduced by Brendan Gregg [@brendangregg](#)
 - KPIs for host based health
- **The Four Golden Signals**
 - Latency, Traffic, Errors, Saturation
 - Covered in the Google SRE Book
 - Extended version of RED
- **RED**
 - Rate, Errors, Duration
 - Introduced by Tom Wilkie [@tom_wilkie](#)
 - KPIs for API based health, SLI focused

Containers?

- Ephemeral
- High Cardinality
- Difficult to Instrument
- Instrument Services, Not Containers



Istio Mixer Provided Telemetry

- Request Count by Response Code
- Request Duration
- Request Size
- Response Size
- Connection Received Bytes
- Connection Sent Bytes
- Connection Duration
- Template Based MetaData (Metric Tags)



RED

- **Rate**

- Requests per second
- First derivative of request count provided by Istio

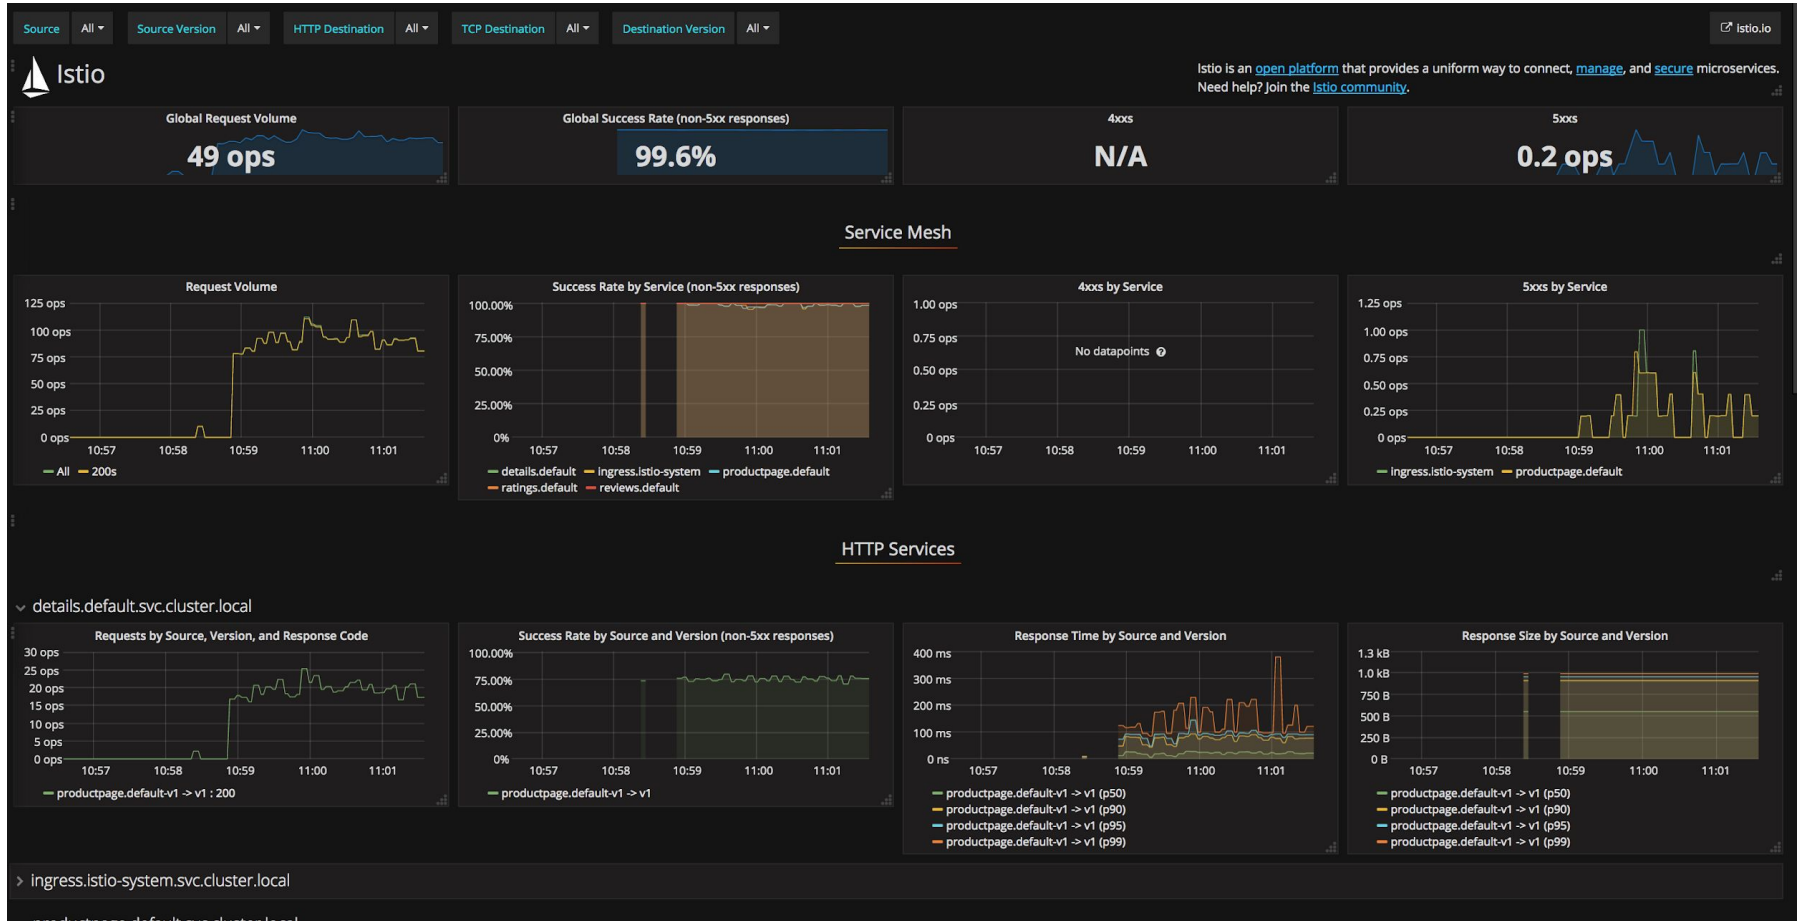
- **Errors**

- Unsuccessful requests per second
- First derivative of failed request count provided by Istio

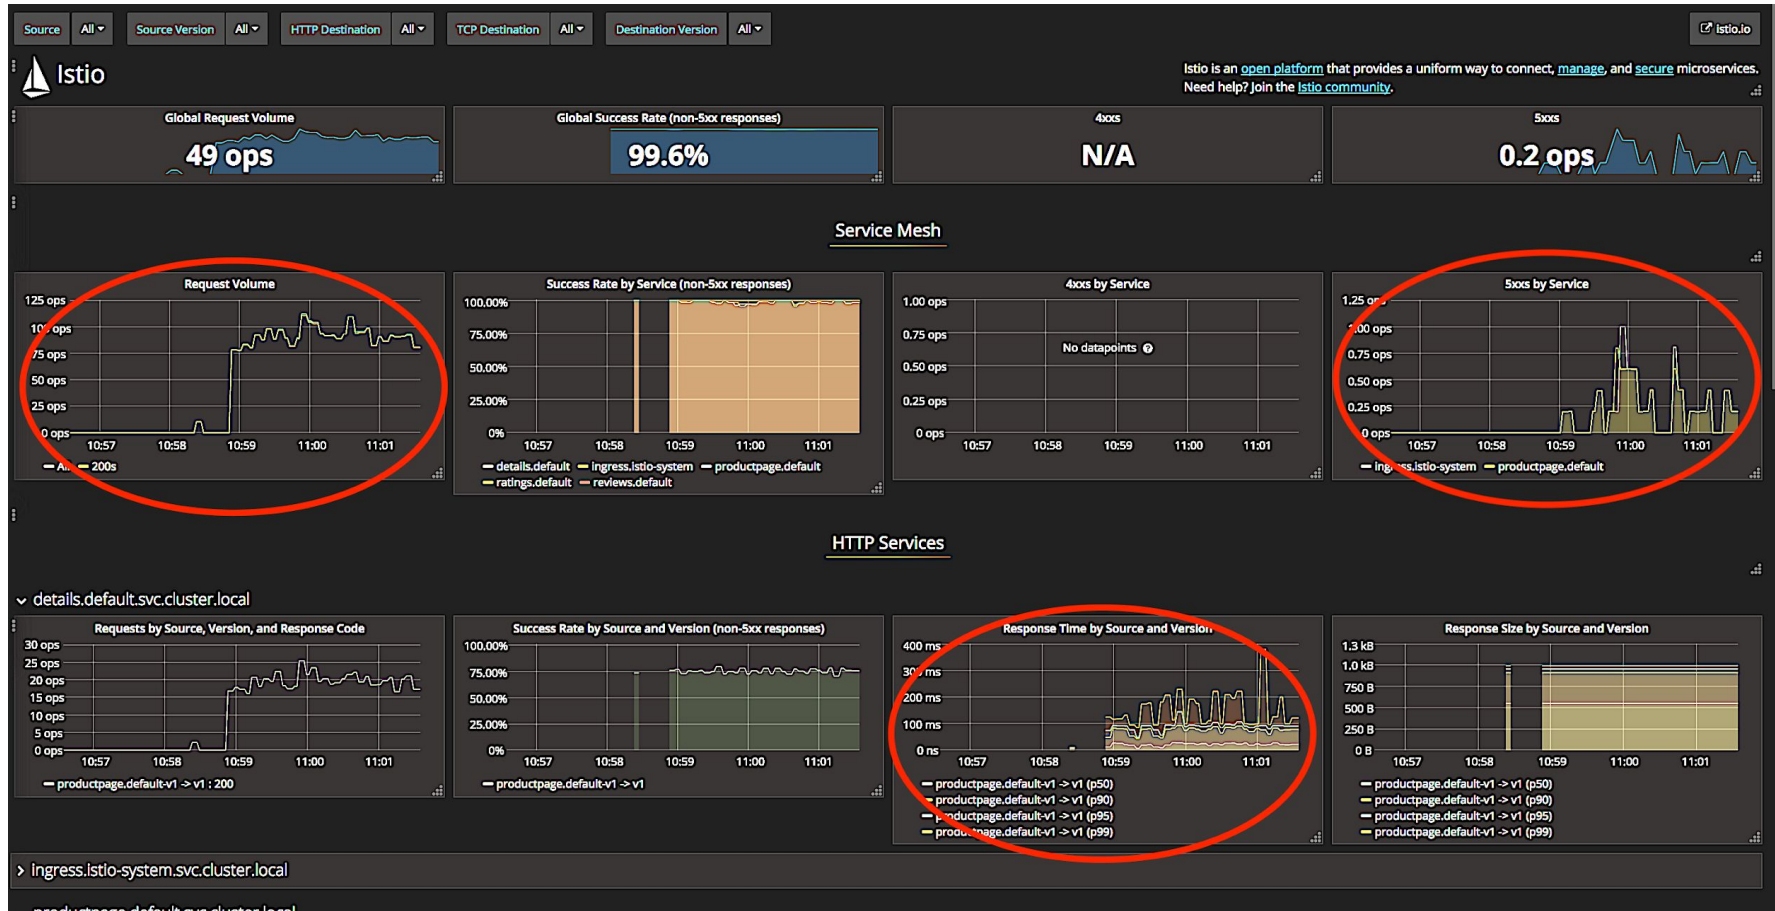
- **Duration**

- Request latency provided by Istio


Istio Grafana Dashboard




Istio Grafana Dashboard




Rate



Errors



Duration




Duration

Problems:

- Percentiles > averages, but have limitations
 - Aggregated metric, fixed time window
 - Cannot be re-aggregated for cluster health
 - Cannot be averaged (common mistake)
- Stored aggregates are outputs, not inputs
- Difficult to measure cluster SLIs
- Leave a lot to be desired for

WE CAN DO BETTER

WE HAVE THE
TECHNOLOGY^W^W
MATH




<https://youtu.be/yCX1Ze3OckKo>

Talk Agenda


- ✓ Istio Overview
- ✓ Service Level Objectives
- ✓ RED Dashboard
- ☐ Histogram Telemetry
- ☐ Istio Metrics Adapter
- ☐ Asking the Right Questions

Histogram




<https://github.com/circonus-labs/circonusllhist>

Log linear histogram




<https://github.com/circonus-labs/circonusllhist>

Duration - Histogram




<https://github.com/circonus-labs/circonusllhist>

Duration - SLI




SLI - “90th percentile latency of requests over past 5 minutes < 1,000ms”


Duration - Modes



Duration - Modes




Duration - Modes




Duration - Heatmap

Load Balancer Request Duration Heatmap ▾



Duration - Heatmap



May 14 2018, 19:10 (1M)

API Requests (histogram) (ms)

0.176k, 2.501k, 2.899k, 2.903k

CAQL 2 (ms)

[16 - 17] 132 of 2903 samples 71% 5% 24%




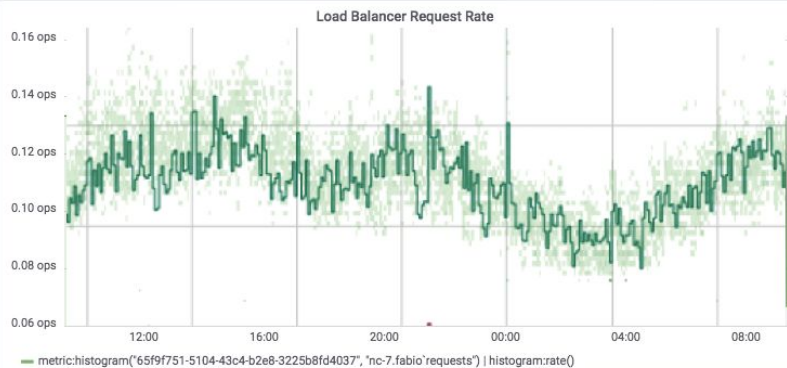
RED

IRONdb RED Dashboard




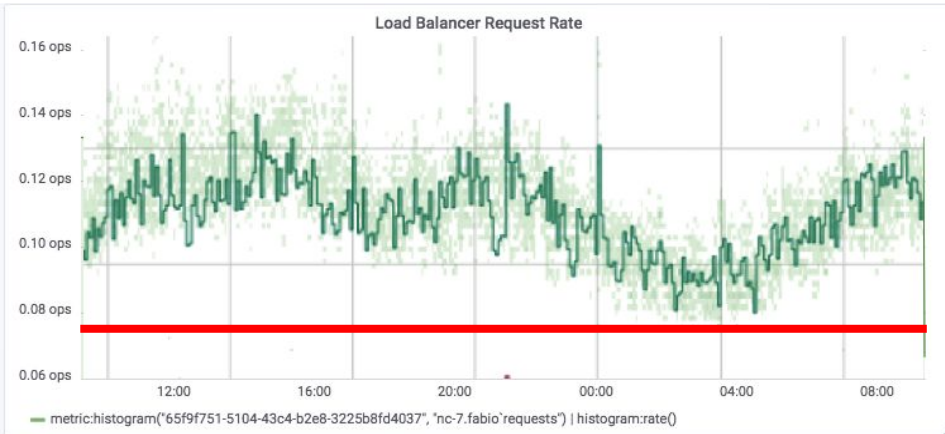


Grafana RED dashboard with IRONdb

This dashboard implements graphs for Rate, Duration, and Error service metrics from the IRONdb datasource.




RED - SLI Alerting



Talk Agenda


- ✓ Istio Overview
- ✓ Service Level Objectives
- ✓ RED Dashboard
- ✓ Histogram Telemetry
- ☐ Istio Metrics Adapter
- ☐ Asking the Right Questions

Istio Metrics Adapter



- Golang based adapter API
- In process (built into the Mixer executable)
 - Out of process for new adapter dev
- Set of handler hooks and YAML files

Istio Metrics Adapter



“SHOW ME THE CODE”

<https://github.com/istio/istio/>

<https://github.com/istio/istio/blob/master/mixer/adapter/circonus>

Istio Metrics Adapter

istio / istio

Unwatch

542

★ Star

8,400

Fork

1,142

Code

Issues 812

Pull requests 110

Wiki

Insights

Branch: master


istio / mixer / adapter / circonus /

Create new file

Upload files





Find file

History

 **guptasu** and **geeknoid** Reference new types from policy/v1beta1 (#5587)

Latest commit 160bf81 3 days ago

..

 config	Reference new types from policy/v1beta1 (#5587)	3 days ago
 operatorconfig	Cleanup attribute bag usage in mixer/pkg/api. (#4228)	3 months ago
 circonus.go	Introduce pkg/ctrlz, Istio's introspection package. (#5123)	a month ago
 circonus_test.go	Circonus metrics adapter (#1737)	6 months ago

Istio Metrics Adapter

📁 bypass	Add a bypass adapter to Mixer for using gRPC backends via inline mode...	a day ago
📁 circonus	Doc updates (#5914)	4 days ago
📁 cloudwatch	Doc updates (#5914)	4 days ago
📁 denier	Doc updates (#5914)	4 days ago
📁 dogstatsd	Doc updates (#5914)	4 days ago
📁 fluentd	Doc updates (#5914)	4 days ago
📁 kubernetesenv	Doc updates (#5914)	4 days ago
📁 list	Doc updates (#5914)	4 days ago
📁 memquota	Doc updates (#5914)	4 days ago
📁 noop	Bump up coverage numbers in a few packages. (#4519)	2 months ago
📁 opa	Doc updates (#5914)	4 days ago
📁 prometheus	Doc updates (#5914)	4 days ago
📁 rbac	Doc updates (#5914)	4 days ago
📁 redisquota	Doc updates (#5914)	4 days ago
📁 servicecontrol	Doc updates (#5914)	4 days ago
📁 solarwinds	Doc updates (#5914)	4 days ago
📁 stackdriver	Doc updates (#5914)	4 days ago
📁 statsd	Doc updates (#5914)	4 days ago
📁 stdio	Doc updates (#5914)	4 days ago



Istio Metrics Adapter

```
// HandleMetric submits metrics to Circonus via circonus-gometrics
func (h *handler) HandleMetric(ctx context.Context, insts
    []*metric.Instance) error {

    for _, inst := range insts {
        metricName := inst.Name
        metricType := h.metrics[metricName]

        switch metricType {

        case config.GAUGE:
            value, _ := inst.Value.(int64)
            h.cm.Gauge(metricName, value)

        case config.COUNTER:
            h.cm.Increment(metricName)
        }
    }
}
```



Istio Metrics Adapter

```
case config.DISTRIBUTION:  
    value, _ := inst.Value.(time.Duration)  
    h.cm.Timing(metricName, float64(value))  
}
```

```
}
```

```
return nil
```

```
}
```

Istio Metrics Adapter

```
handler struct {  
    cm      *cgm.CirconusMetrics  
    env     adapter.Env  
    metrics map[string]config.Params_MetricInfo_Type  
    cancel  context.CancelFunc  
}
```

And some YAML

metrics:

- name: requestcount.metric.istio-system
type: COUNTER
- name: requestduration.metric.istio-system
type: DISTRIBUTION
- name: requestsize.metric.istio-system
type: GAUGE
- name: responsesize.metric.istio-system
type: GAUGE

Buffer metrics, then report

```
env.ScheduleDaemon(  
    func() {  
  
        ticker := time.NewTicker(b.adpCfg.SubmissionInterval)  
        for {  
            select {  
            case <-ticker.C:  
                cm.Flush()  
            case <-adapterContext.Done():  
                ticker.Stop()  
                cm.Flush()  
                return  
            }  
        }  
    }  
})
```


Talk Agenda

- ✓ Istio Overview
- ✓ Service Level Objectives
- ✓ RED Dashboard
- ✓ Histogram Telemetry
- ✓ Istio Metrics Adapter
- ☐ Asking the Right Questions

Your boss wants to know

- How many users got angry on the Tuesday slowdown after the big marketing promotion?
- Are we over-provisioned or under-provisioned on our purchasing checkout service?
- Other business centric questions



The Slowdown

- Marketing launched a new product
- Users complained the site was slow
- Median human reaction time is 215 ms [1]
- If users get angry (rage clicks) when requests take more than 500 ms, how many users got angry?


[1] - <https://www.humanbenchmark.com/tests/reactiontime/>

The Slowdown

1. Record all service request latencies as distribution
2. Plot as a heatmap
3. Calculate percentage of requests that exceed 500ms SLI using inverse percentiles
4. Multiply result by total number requests, integrate over time

The Slowdown

4 million slow requests




L R Cumulative number of requests exceeded the 500 millisecond SLO ⊙ × ▾

L R circonus-demo circonus.net json`GET`/getState (on demo-replay.circonus.net, from Chicago, IL, US) ⊙ × ▾

Under or Over Provisioned?

- “It depends”
- Time of day, day of week
- Special events
- Behavior under load
- Latency bands shed some light

Latency Bands



		⚙️ <i>May 23 2018, 18:26 (1M)</i>
L R 0	Count below 25ms	2.065k
L R 0	Count 25-100ms	0.363k
L R 0	Count 100-250 ms	0.019k
L R 0	Count 250ms - 1,000ms	0.042k
L R 0	Count above 1,000ms	0.023k

Conclusions

- Monitor services, not containers
- Record distributions, not aggregates
- Istio gives you RED metrics for free
- Use math to ask the right questions

Thank you! Questions?

Tweet me
[@phredmoyer](https://twitter.com/phredmoyer)

