SPRING CLOUD + KUBERNETES + ISTIO = ?

A MACRO PERSPECTIVE ON THE TOOLBOX FOR MICROSERVICES

MAGNUS LARSSON

CADEC 2020.01.23 & 2020.01.29 | CALLISTAENTERPRISE.SE

CALLISTA

AGENDA

- Why?
- Challenges
- Open Source to the rescue!
- Overlaps
- Demo
- Summary

WHY?

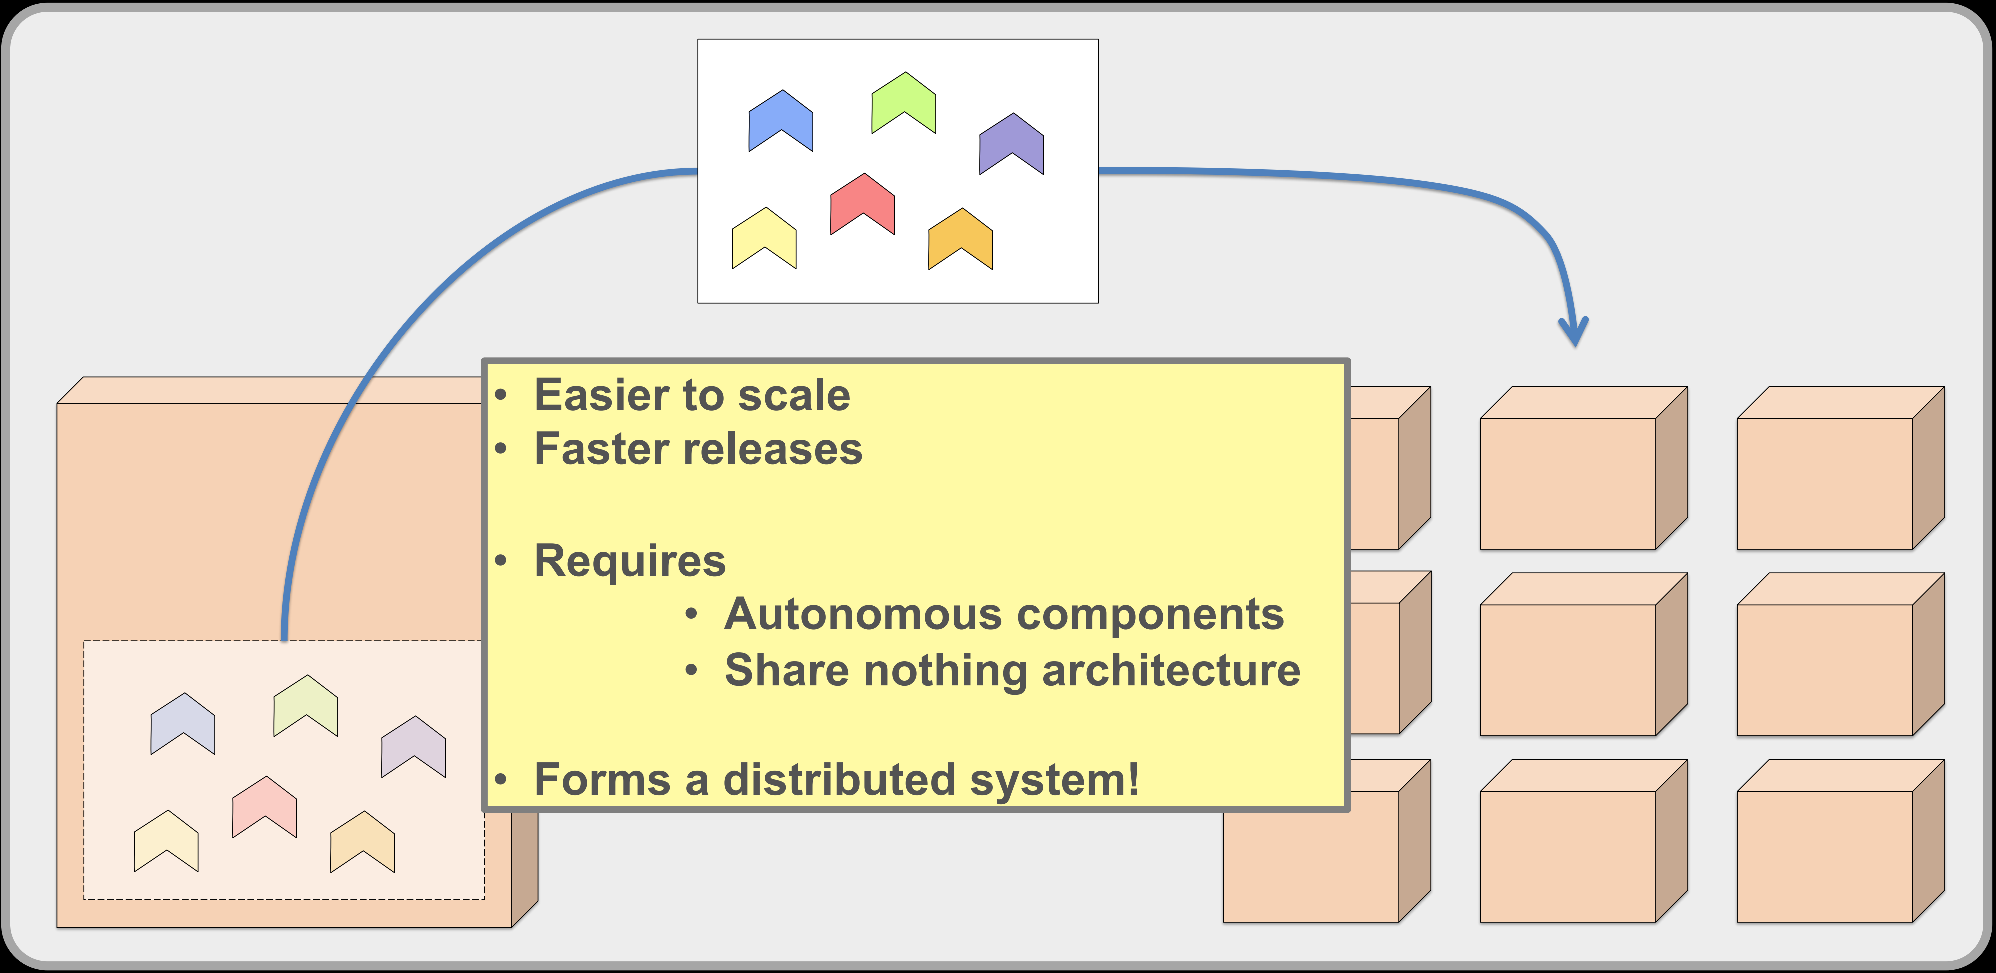

- Easier to scale
- Faster releases

- Requires
  - Autonomous components
  - Share nothing architecture

- Forms a distributed system!

CHALLENGES

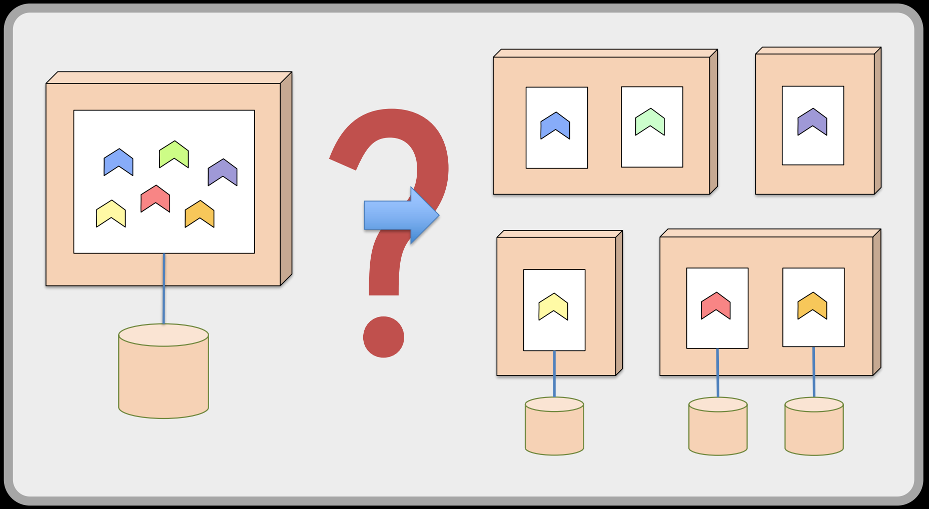

CALLISTA

CHALLENGES

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

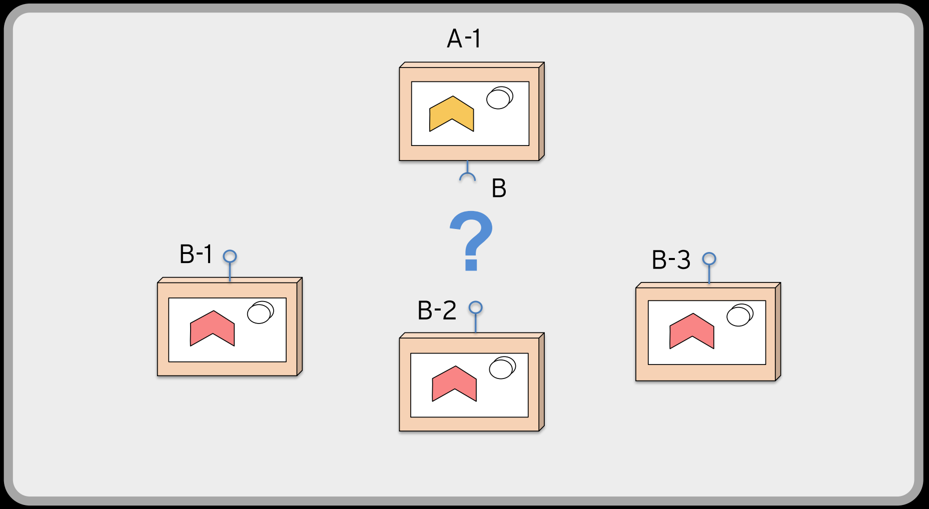

CHALLENGES

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?

HOW TO PROTECT PUBLIC SERVICES?

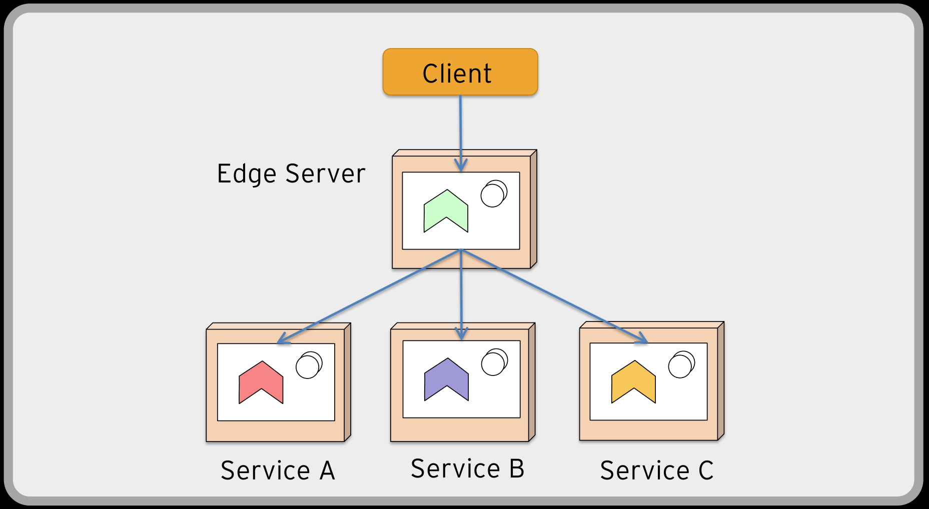

CHALLENGES

DISCOVERY SERVER
WHERE ARE THE SERVICES?
WHICH SERVICE TO CALL?

EDGE SERVER
HOW TO HIDE PRIVATE SERVICES?
HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION
WHERE IS MY CONFIGURATION?
ARE ALL SERVICES
CONFIGURATION UP TO DATE?

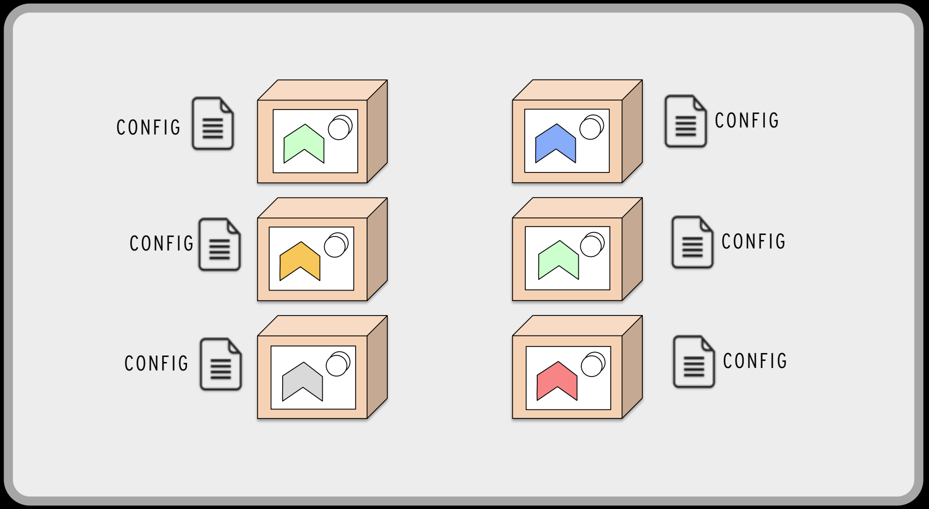

CALLISTA

CHALLENGES

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?

HOW TO PROTECT PUBLIC SERVICES?

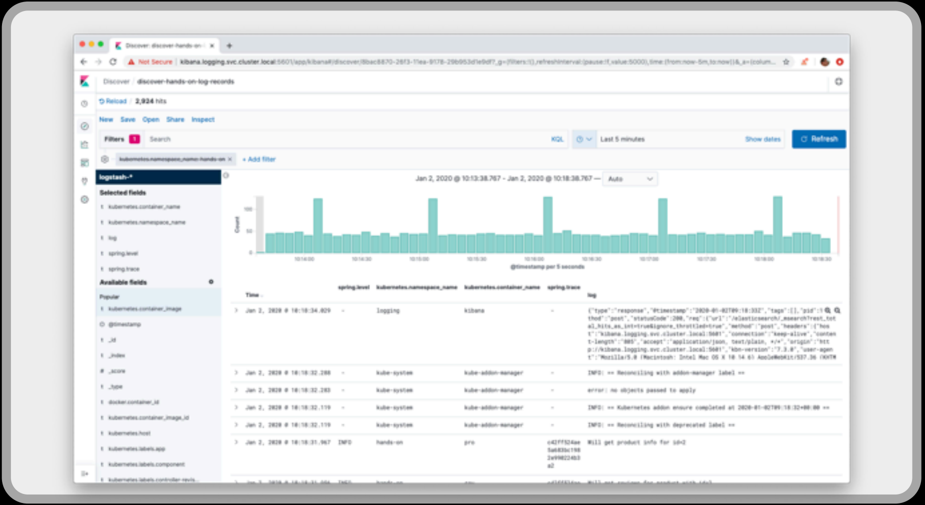
Discover: discover-hands-on-l
Not Secure
kibana.logging.svc.cluster.local:5601/app/kibana#/discover/8bac8870-2613-11ea-9178-29b953d1e9df?_g=(filters:!(),refreshInterval:(pause:!f,value:5000),time:(from:now-5m,to:now))&_a=(colum...
Discover / discover-hands-on-log-records
Reload / 2,924 hits
New
Save
Open
Share
Inspect
Filters 1
Search
KQL
Last 5 minutes
Show dates
Refresh
kubernetes.namespace_name: hands-on
+ Add filter
logstash-*
Selected fields
kubernetes.container_name
kubernetes.namespace_name
log
spring.level
spring.trace
Available fields
Popular
kubernetes.container_image
@timestamp
_id
_index
_score
_type
docker.container_id
kubernetes.container_image_id
kubernetes.host
kubernetes.labels.app
kubernetes.labels.component
kubernetes.labels.controller-revis...
Jan 2, 2020 @ 10:13:38.767 - Jan 2, 2020 @ 10:18:38.767 —
Auto
Count
100
50
0
10:14:00
10:14:30
10:15:00
10:15:30
10:16:00
10:16:30
10:17:00
10:17:30
10:18:00
10:18:30
@timestamp per 5 seconds

Time	spring.level	kubernetes.namespace_name	kubernetes.container_name	spring.trace	log
Jan 2, 2020 @ 10:18:34.029	-	logging	kibana	-	{"type":"response","@timestamp":"2020-01-02T09:18:33Z","tags":[],"pid":1,"method":"post","statusCode":200,"req":{"url":"/elasticsearch/_msearch?rest_total_hits_as_int=true&ignore_throttled=true","method":"post","headers":{"host":"kibana.logging.svc.cluster.local:5601","connection":"keep-alive","content-length":"885","accept":"application/json, text/plain, */*","origin":"http://kibana.logging.svc.cluster.local:5601","kbn-version":"7.3.0","user-agent":"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_14_6) AppleWebKit/537.36 (KHTM
Jan 2, 2020 @ 10:18:32.288	-	kube-system	kube-addon-manager	-	INFO: == Reconciling with addon-manager label ==
Jan 2, 2020 @ 10:18:32.283	-	kube-system	kube-addon-manager	-	error: no objects passed to apply
Jan 2, 2020 @ 10:18:32.119	-	kube-system	kube-addon-manager	-	INFO: == Kubernetes addon ensure completed at 2020-01-02T09:18:32+00:00 ==
Jan 2, 2020 @ 10:18:32.119	-	kube-system	kube-addon-manager	-	INFO: == Reconciling with deprecated label ==
Jan 2, 2020 @ 10:18:31.967	INFO	hands-on	pro	c42ff524ae5a683bc1982e998224b3a2	Will get product info for id=2

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?

ARE ALL SERVICES

CONFIGURATION UP TO DATE?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS

FROM DIFFERENT SERVICES?

CHALLENGES

DISCOVERY SERVER

WHERE ARE THE SERVICES?
WHICH SERVICE TO CALL?

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?
HOW TO PROTECT PUBLIC SERVICES?

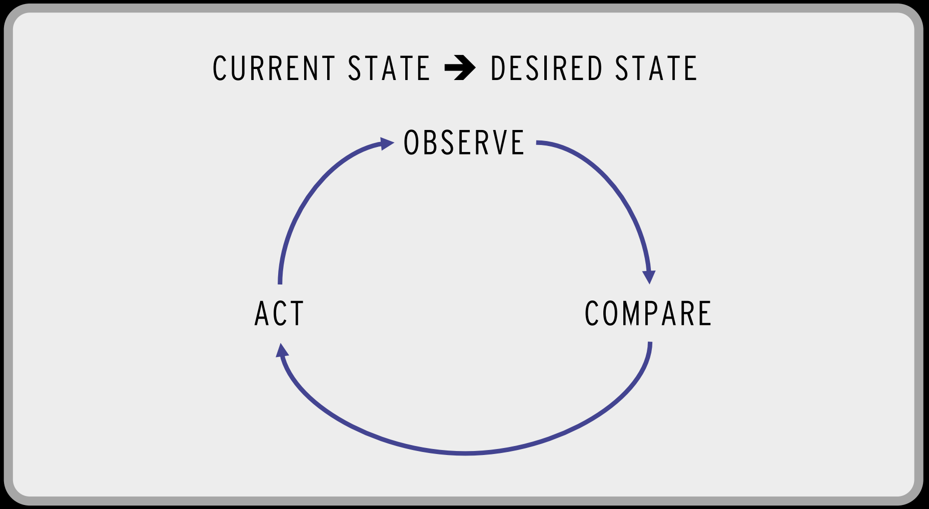

CURRENT STATE → DESIRED STATE

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?
ARE ALL SERVICES
CONFIGURATION UP TO DATE?

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS
FROM DIFFERENT SERVICES?

CHALLENGES

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?

HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?

ARE ALL SERVICES CONFIGURATION UP TO DATE?

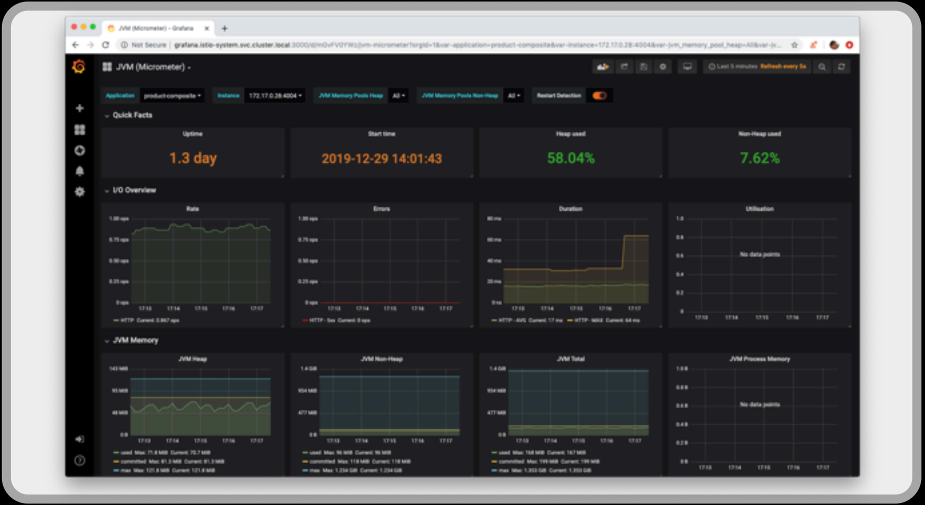

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS FROM DIFFERENT SERVICES?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CHALLENGES

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?

HOW TO PROTECT PUBLIC SERVICES?

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?

ARE ALL SERVICES CONFIGURATION UP TO DATE?

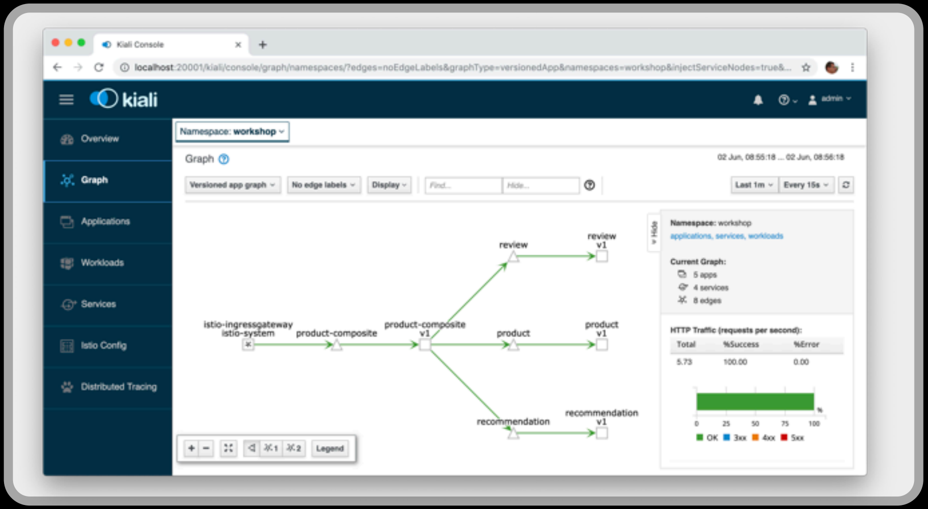

Kiali Console

localhost:20001/kiali/console/graph/namespaces/?edges=noEdgeLabels&graphType=versionedApp&namespaces=workshop&injectServiceNodes=true&...

Namespace: workshop

02 Jun, 08:55:18 ... 02 Jun, 08:56:18

Versioned app graph · No edge labels · Display · Last 1m · Every 15s

Namespace: workshop

applications, services, workloads

Current Graph:
- 5 apps
- 4 services
- 8 edges

HTTP Traffic (requests per second):

Total	%Success	%Error
5.73	100.00	0.00

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS FROM DIFFERENT SERVICES?

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CHALLENGES

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?
HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?
ARE ALL SERVICES
CONFIGURATION UP TO DATE?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS
FROM DIFFERENT SERVICES?

DISCOVERY SERVER

WHERE ARE THE SERVICES?
WHICH SERVICE TO CALL?

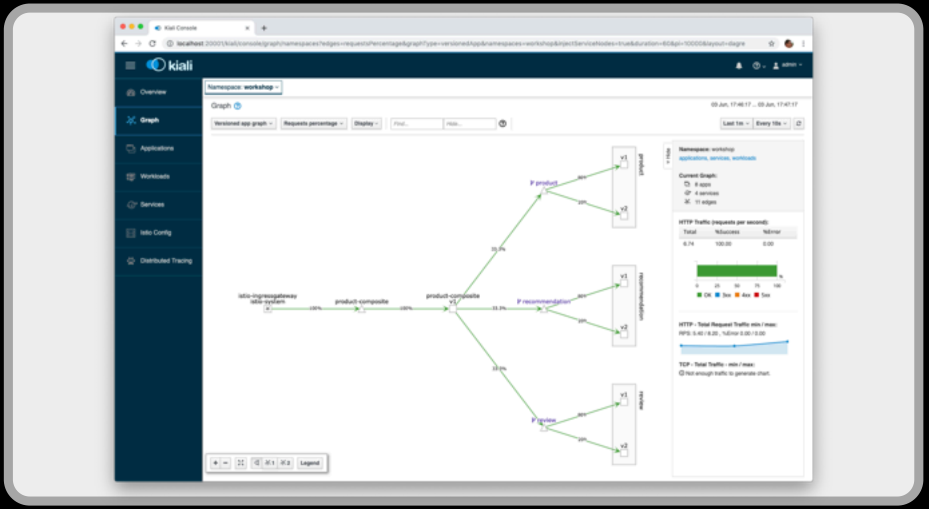

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

TRAFFIC MANAGMENT

HOW TO CONTROL ROUTING?

- RATE LIMITING
- CANARY & BLUE/GREEN UPGRADES

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CHALLENGES

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?

HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?

ARE ALL SERVICES

CONFIGURATION UP TO DATE?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS

FROM DIFFERENT SERVICES?

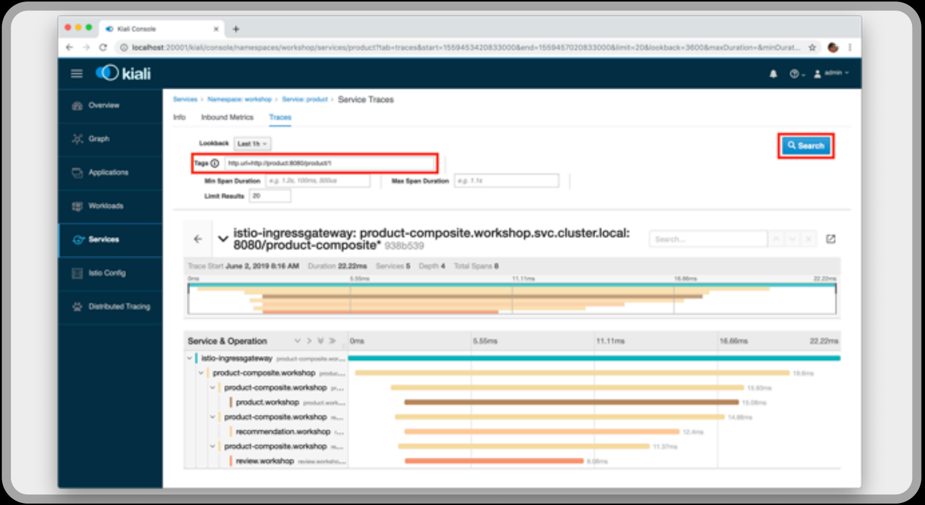

localhost:20001/kiali/console/namespaces/workshop/services/product?tab=traces&start=1559453420833000&end=1559457020833000&limit=20&lookback=3600&maxDuration=&minDurat...

Services › Namespace: workshop › Service: product › Service Traces

Lookback: Last 1h

Tags: http.url=http://product:8080/product/1

Min Span Duration: e.g. 1.2s, 100ms, 500us

Max Span Duration: e.g. 1.1s

Limit Results: 20

istio-ingressgateway: product-composite.workshop.svc.cluster.local:8080/product-composite* 938b539

Trace Start June 2, 2019 8:16 AM Duration 22.22ms Services 5 Depth 4 Total Spans 8

DISTRIBUTED TRACING

WHO IS CALLING WHO?

TRAFFIC MANAGMENT

HOW TO CONTROL ROUTING?

- RATE LIMITING
- CANARY & BLUE/GREEN UPGRADES

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

CHALLENGES

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?
HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?
ARE ALL SERVICES CONFIGURATION UP TO DATE?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS FROM DIFFERENT SERVICES?

DISCOVERY SERVER

WHERE ARE THE SERVICES?
WHICH SERVICE TO CALL?

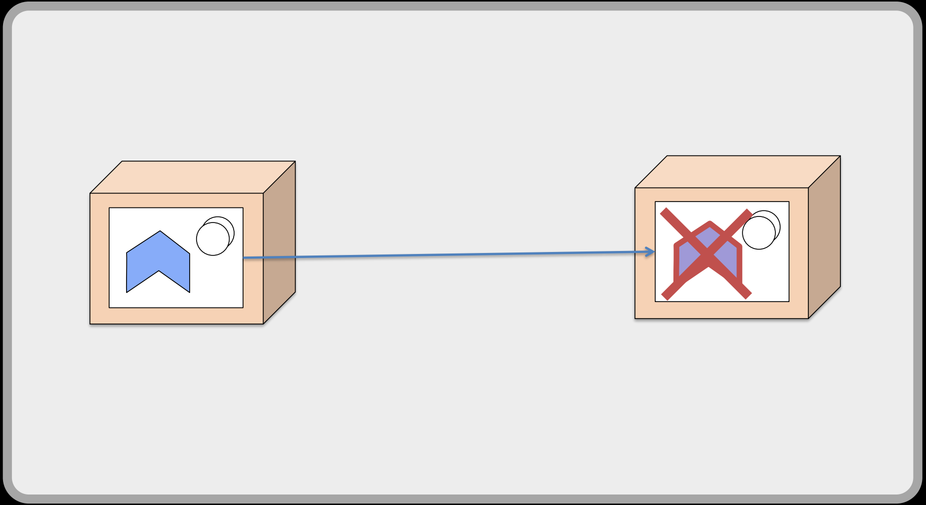

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

RESILIENCE

HOW TO HANDLE FAULTS?

- SLOW OR NO RESPONSE
- TEMPORARY FAULTS
- OVERLOAD

DISTRIBUTED TRACING

WHO IS CALLING WHO?

TRAFFIC MANAGMENT

HOW TO CONTROL ROUTING?

- RATE LIMITING
- CANARY & BLUE/GREEN UPGRADES

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CALLISTA

REQUIRED CAPABILITIES!

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?

HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?

ARE ALL SERVICES CONFIGURATION UP TO DATE?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS FROM DIFFERENT SERVICES?

DISCOVERY SERVER

WHERE ARE THE SERVICES?

WHICH SERVICE TO CALL?

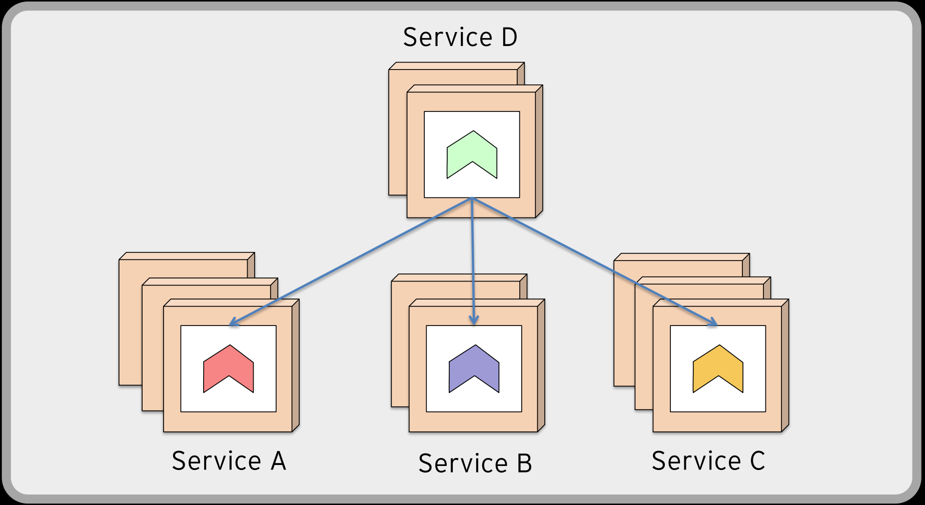

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

RESILIENCE

HOW TO HANDLE FAULTS?

- SLOW OR NO RESPONSE
- TEMPORARY FAULTS
- OVERLOAD

DISTRIBUTED TRACING

WHO IS CALLING WHO?

TRAFFIC MANAGMENT

HOW TO CONTROL ROUTING?

- RATE LIMITING
- CANARY & BLUE/GREEN UPGRADES

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CALLISTA

WHERE ARE WE?

- Why?
- Challenges
- **Open Source to the rescue!**
- Overlaps
- Demo
- Summary

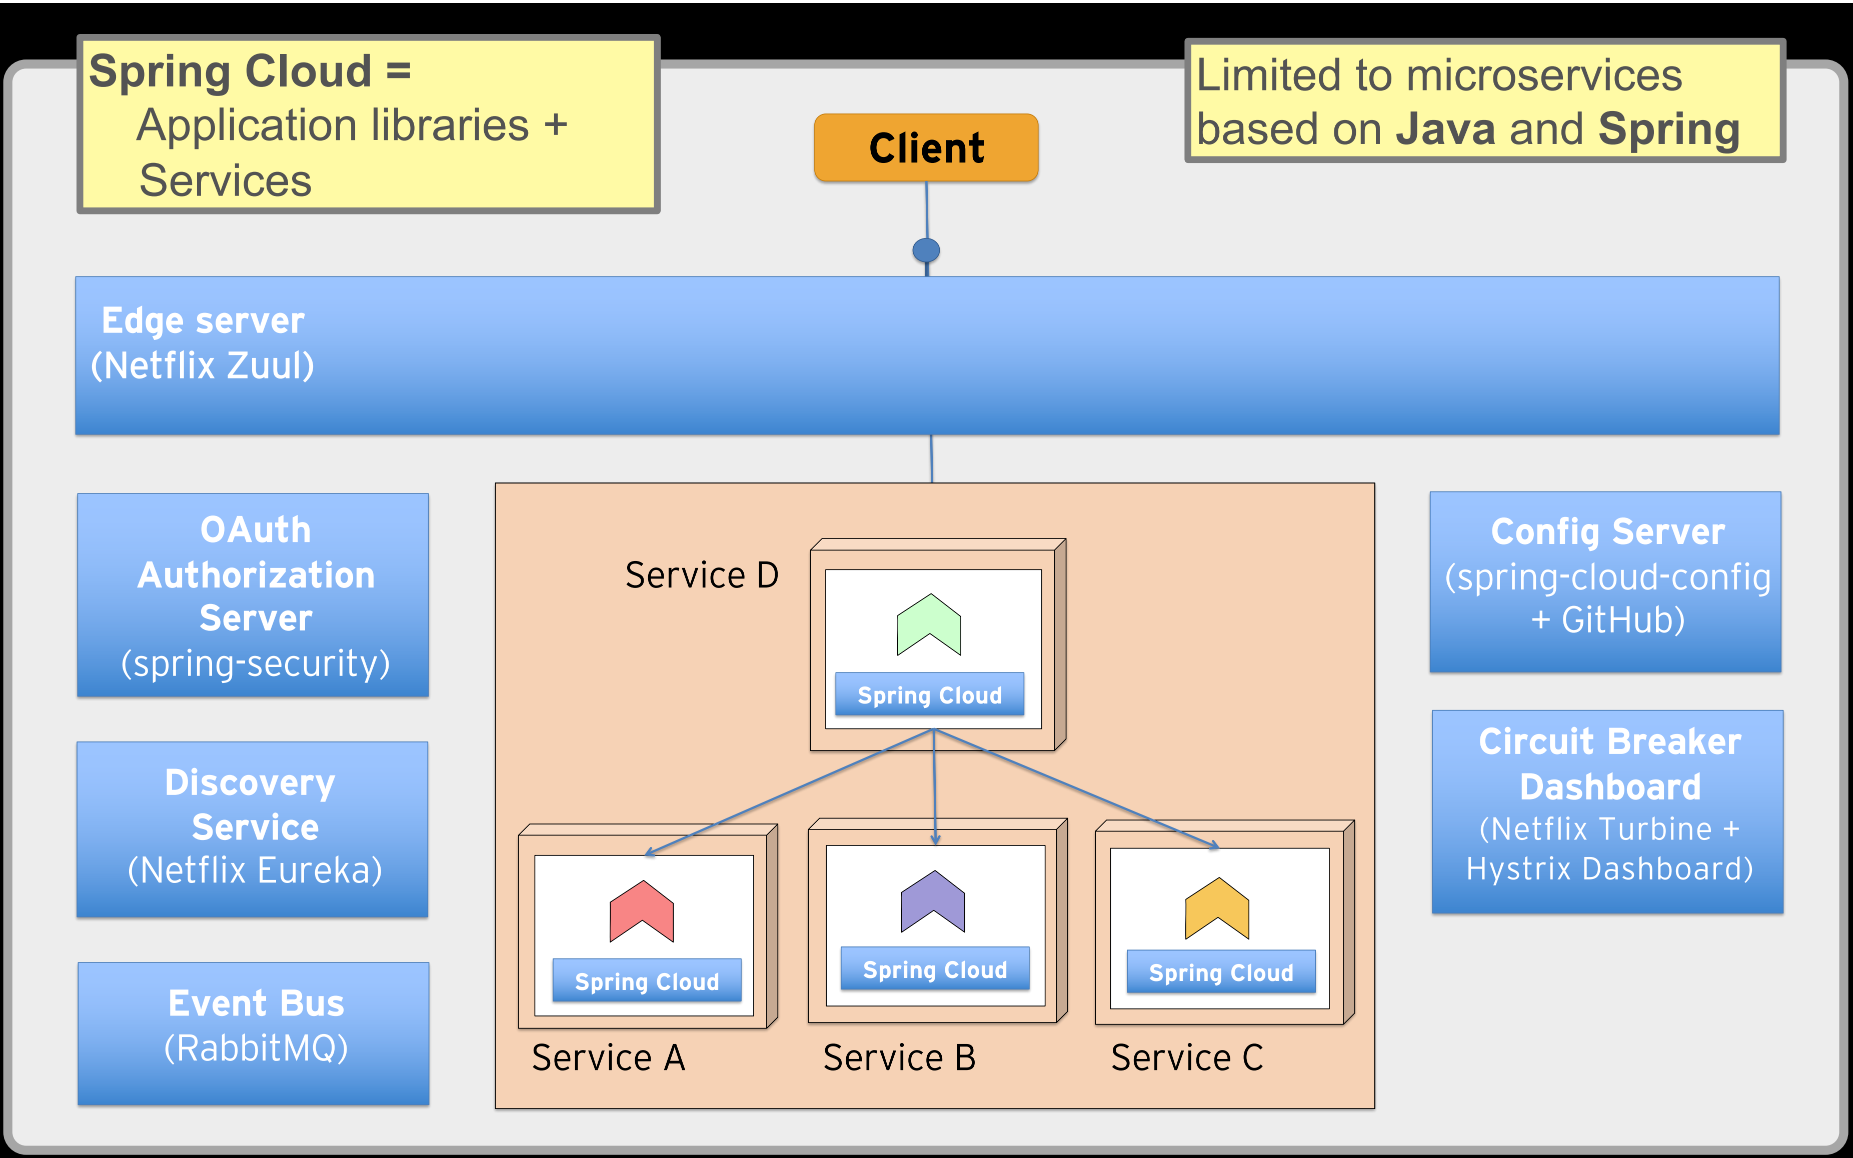

Spring Cloud =
Application libraries +
Services

Limited to microservices based on **Java** and **Spring**

THE EVOLUTION

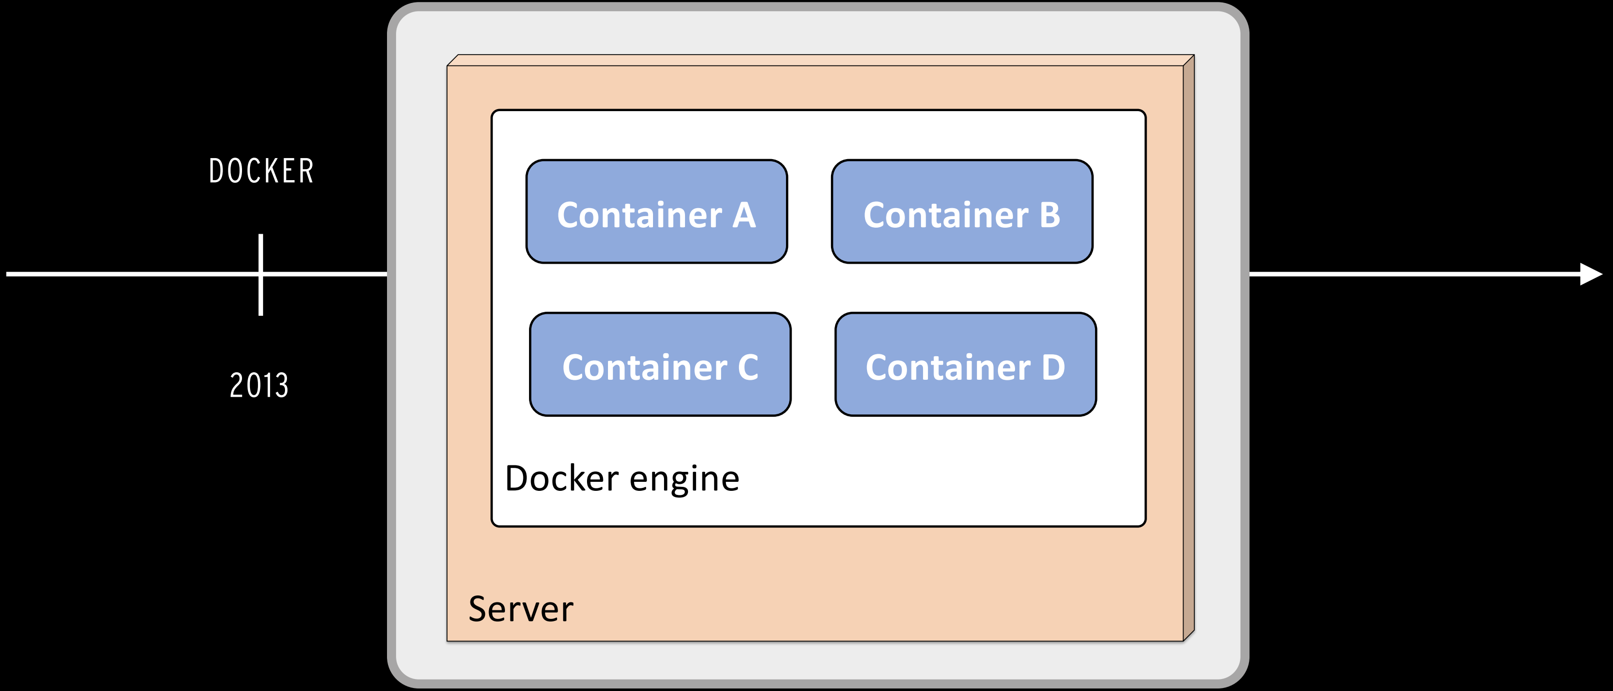

CALLISTA

THE EVOLUTION

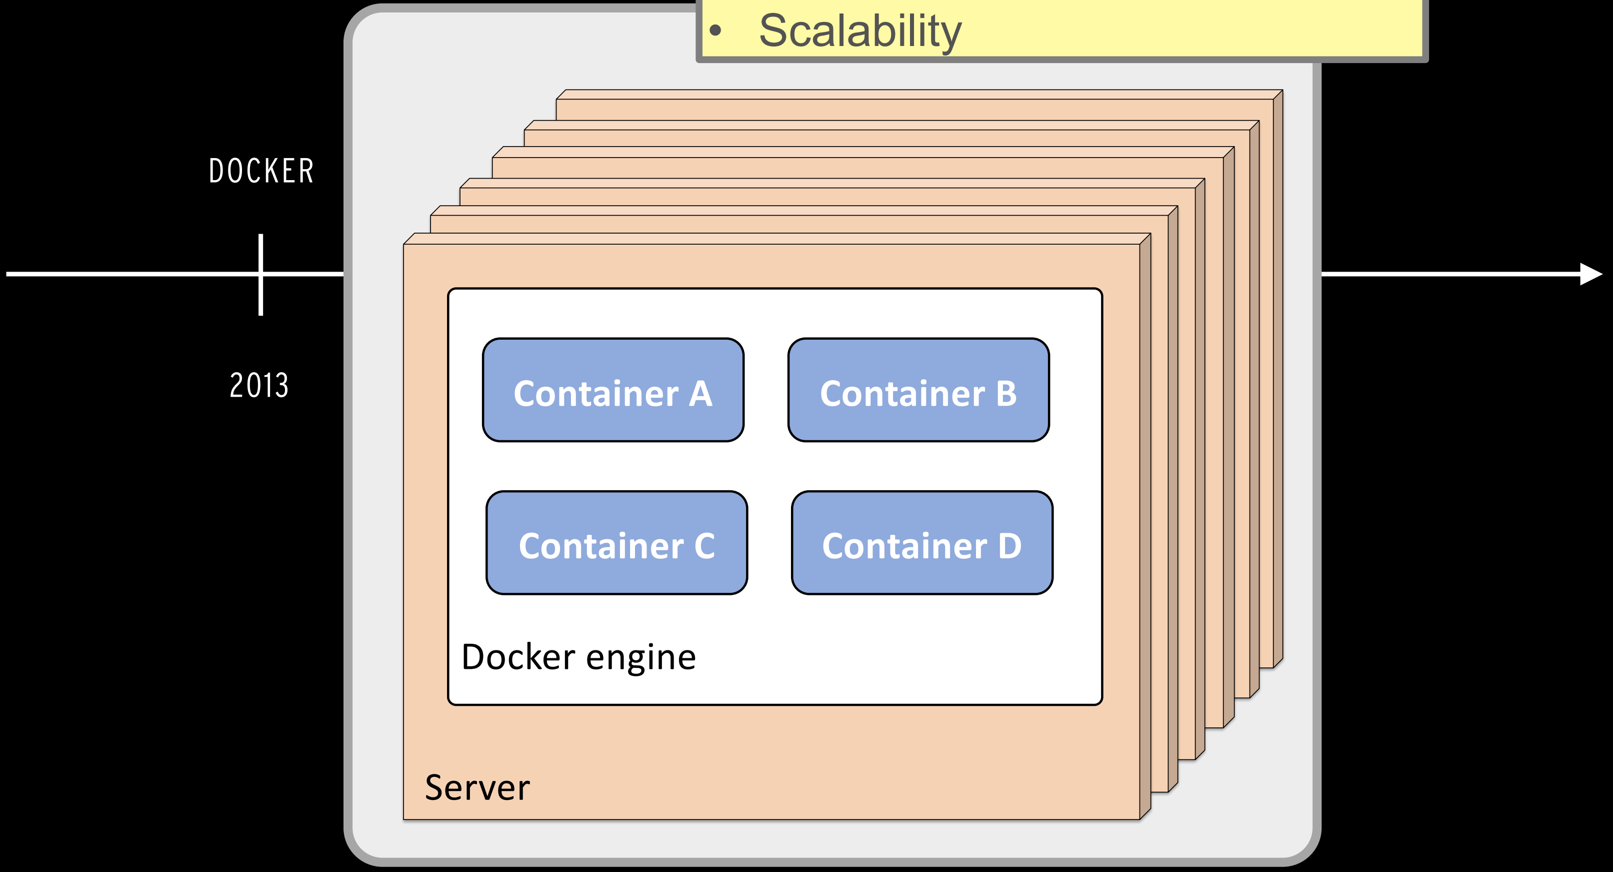

>1 server required in production

- High Availability
- Scalability

CALLISTA

THE EVOLUTION

Kubernetes: A Container Orchestrator

- A cluster of servers running Docker engine acting as one big server
- Enforces *actual state* = *desired state*

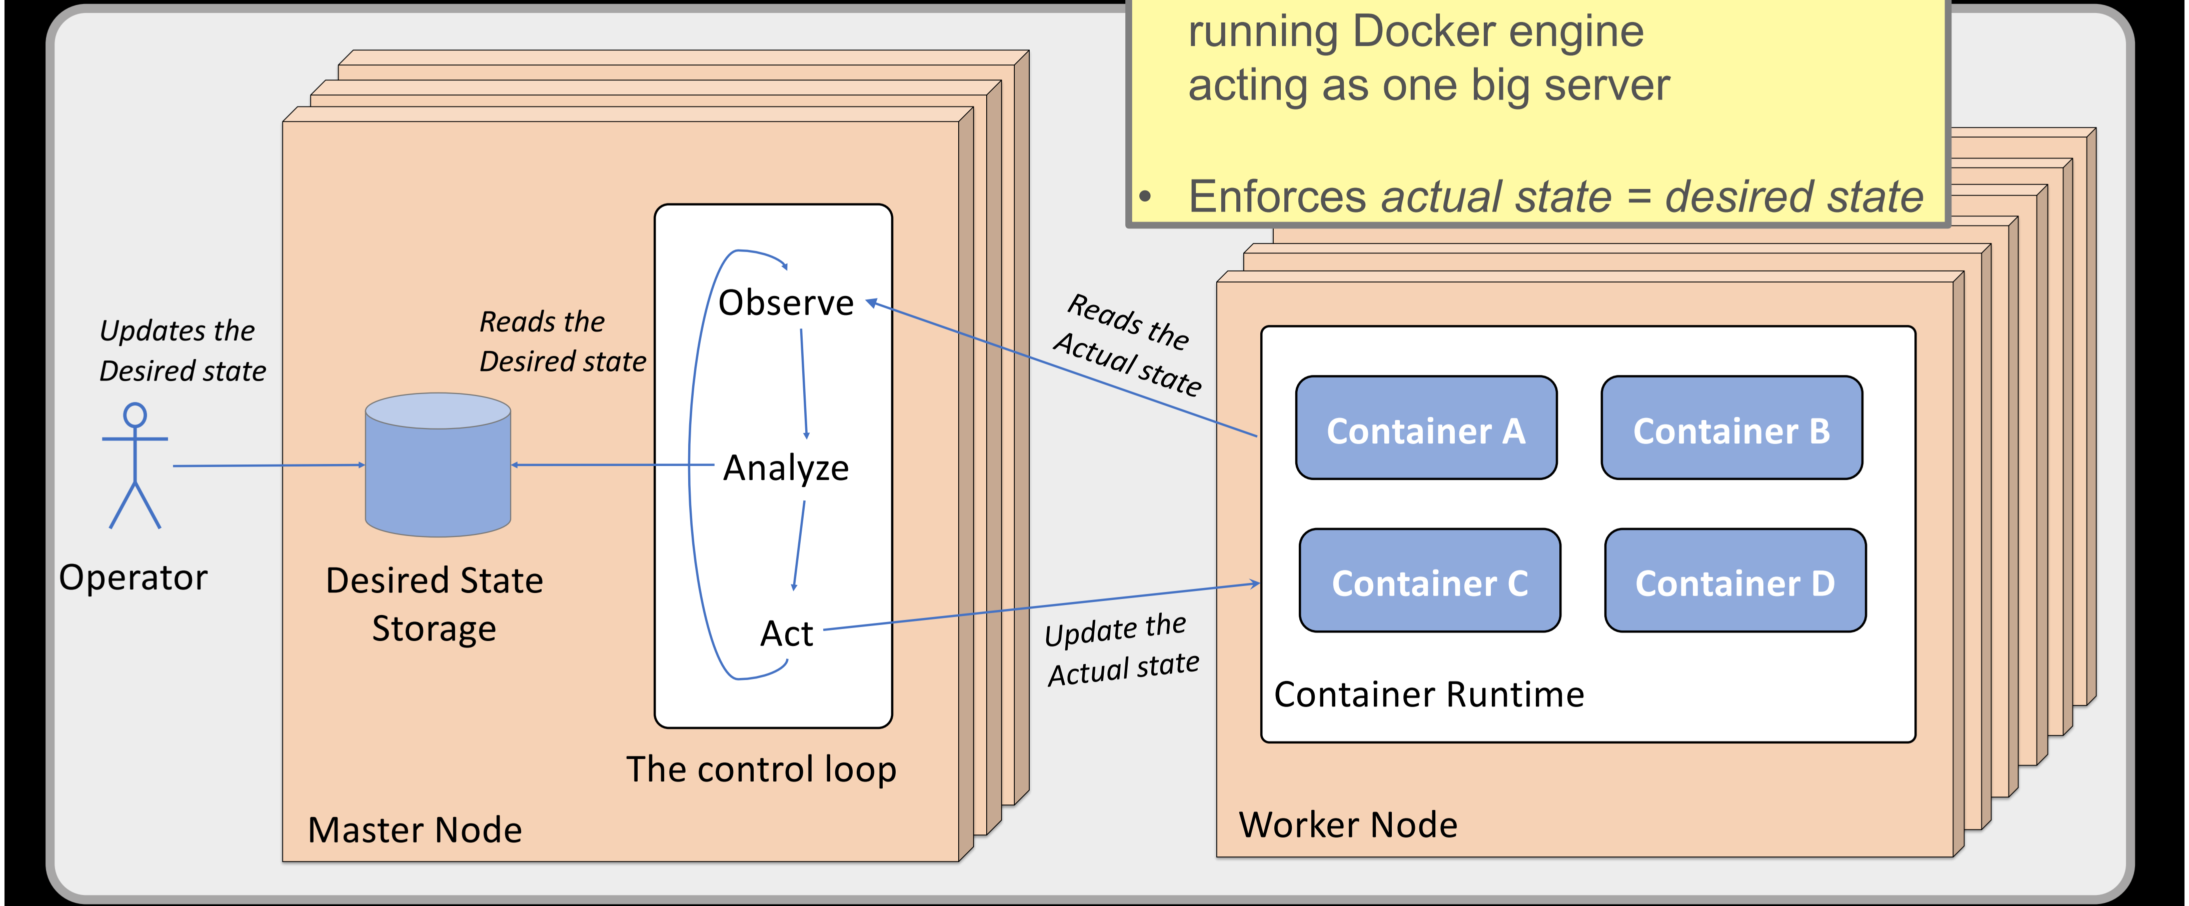

CALLISTA

THE EVOLUTION

ISTIO: Service mesh

Observability, Security, Resilience and Traffic Management

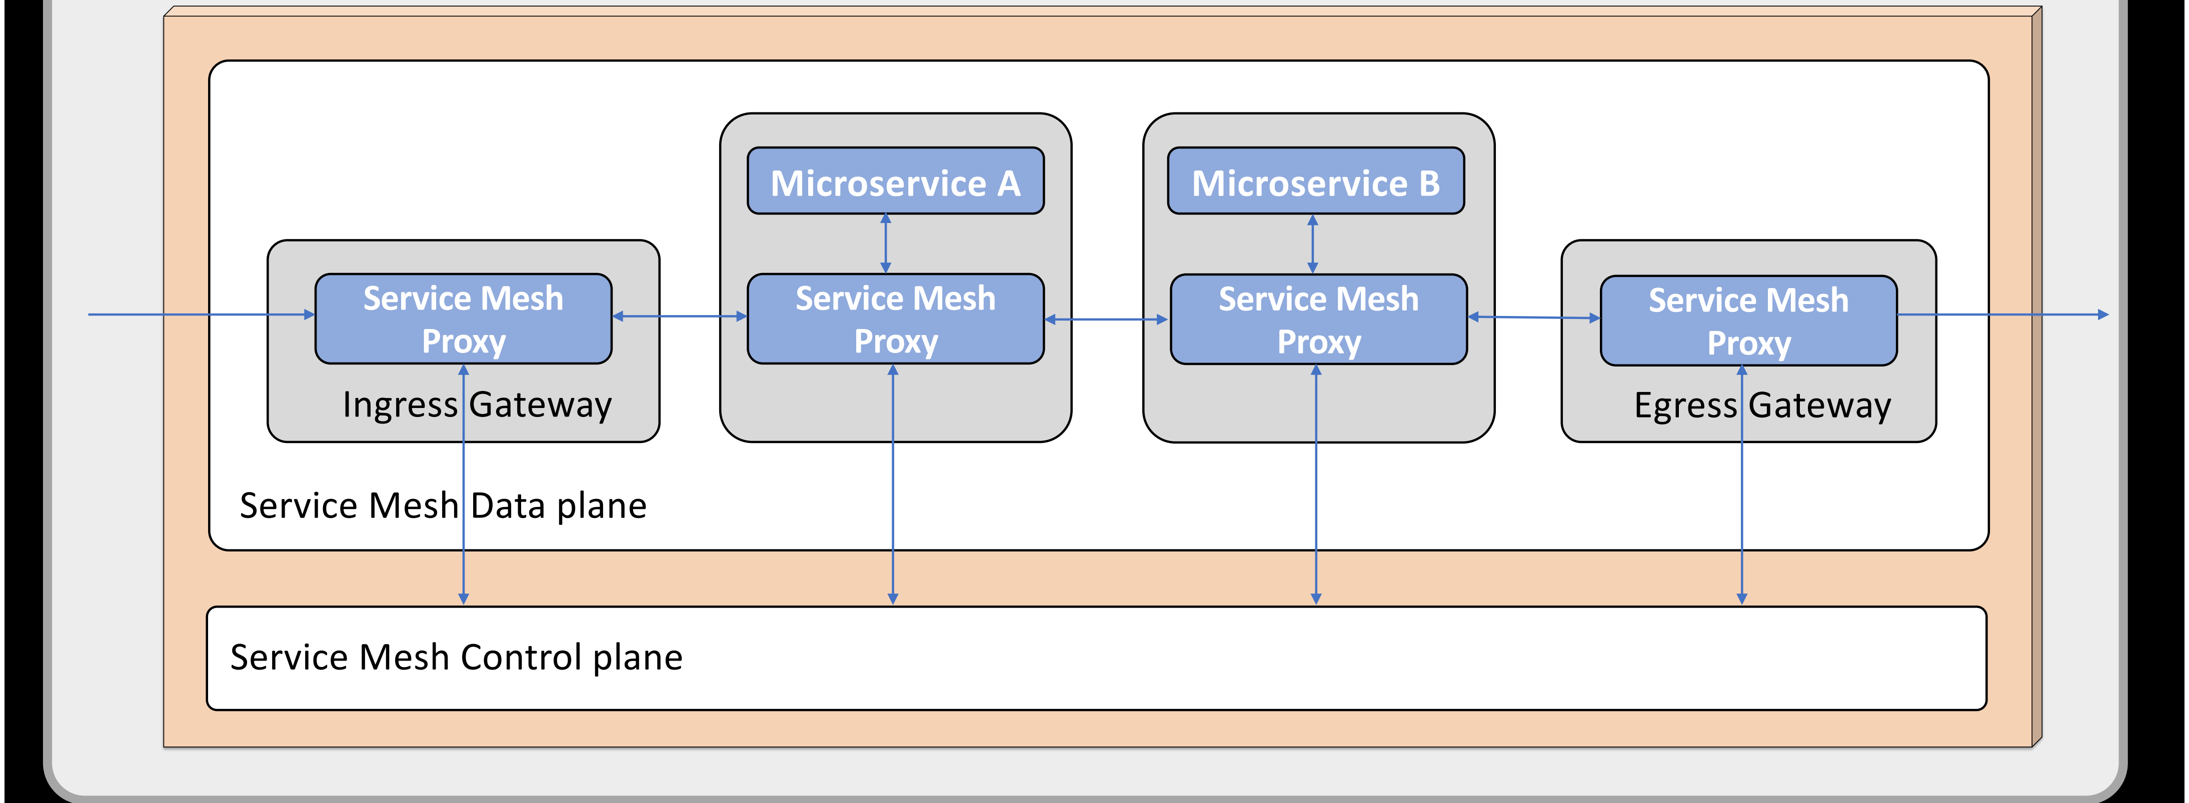

THE EVOLUTION

100% OPEN SOURCE

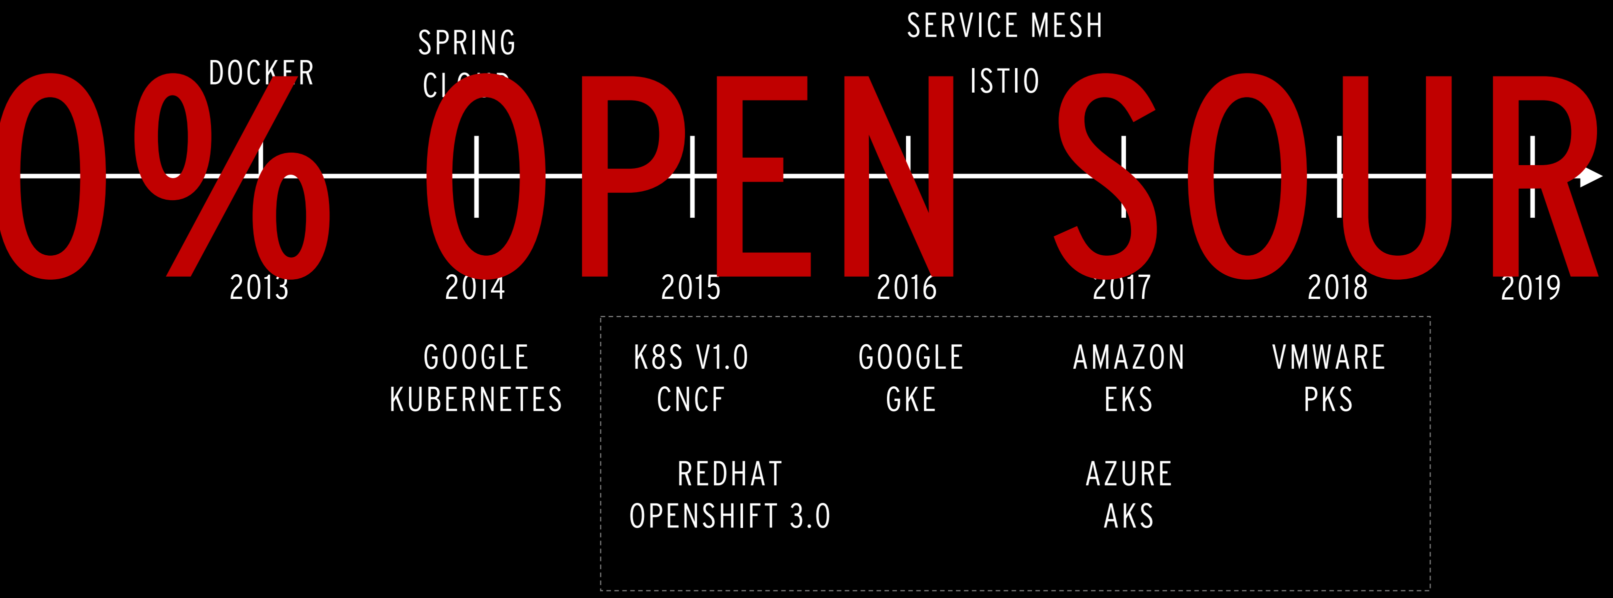

DOCKER
2013

SPRING CLOUD
2014
GOOGLE KUBERNETES

2015
K8S V1.0 CNCF
REDHAT OPENSHIFT 3.0

SERVICE MESH
ISTIO
2016
GOOGLE GKE

2017
AMAZON EKS
AZURE AKS

2018
VMWARE PKS

2019

CAPABILITY MAPPING

SPRING CLOUD

KUBERNETES

ISTIO

EFK

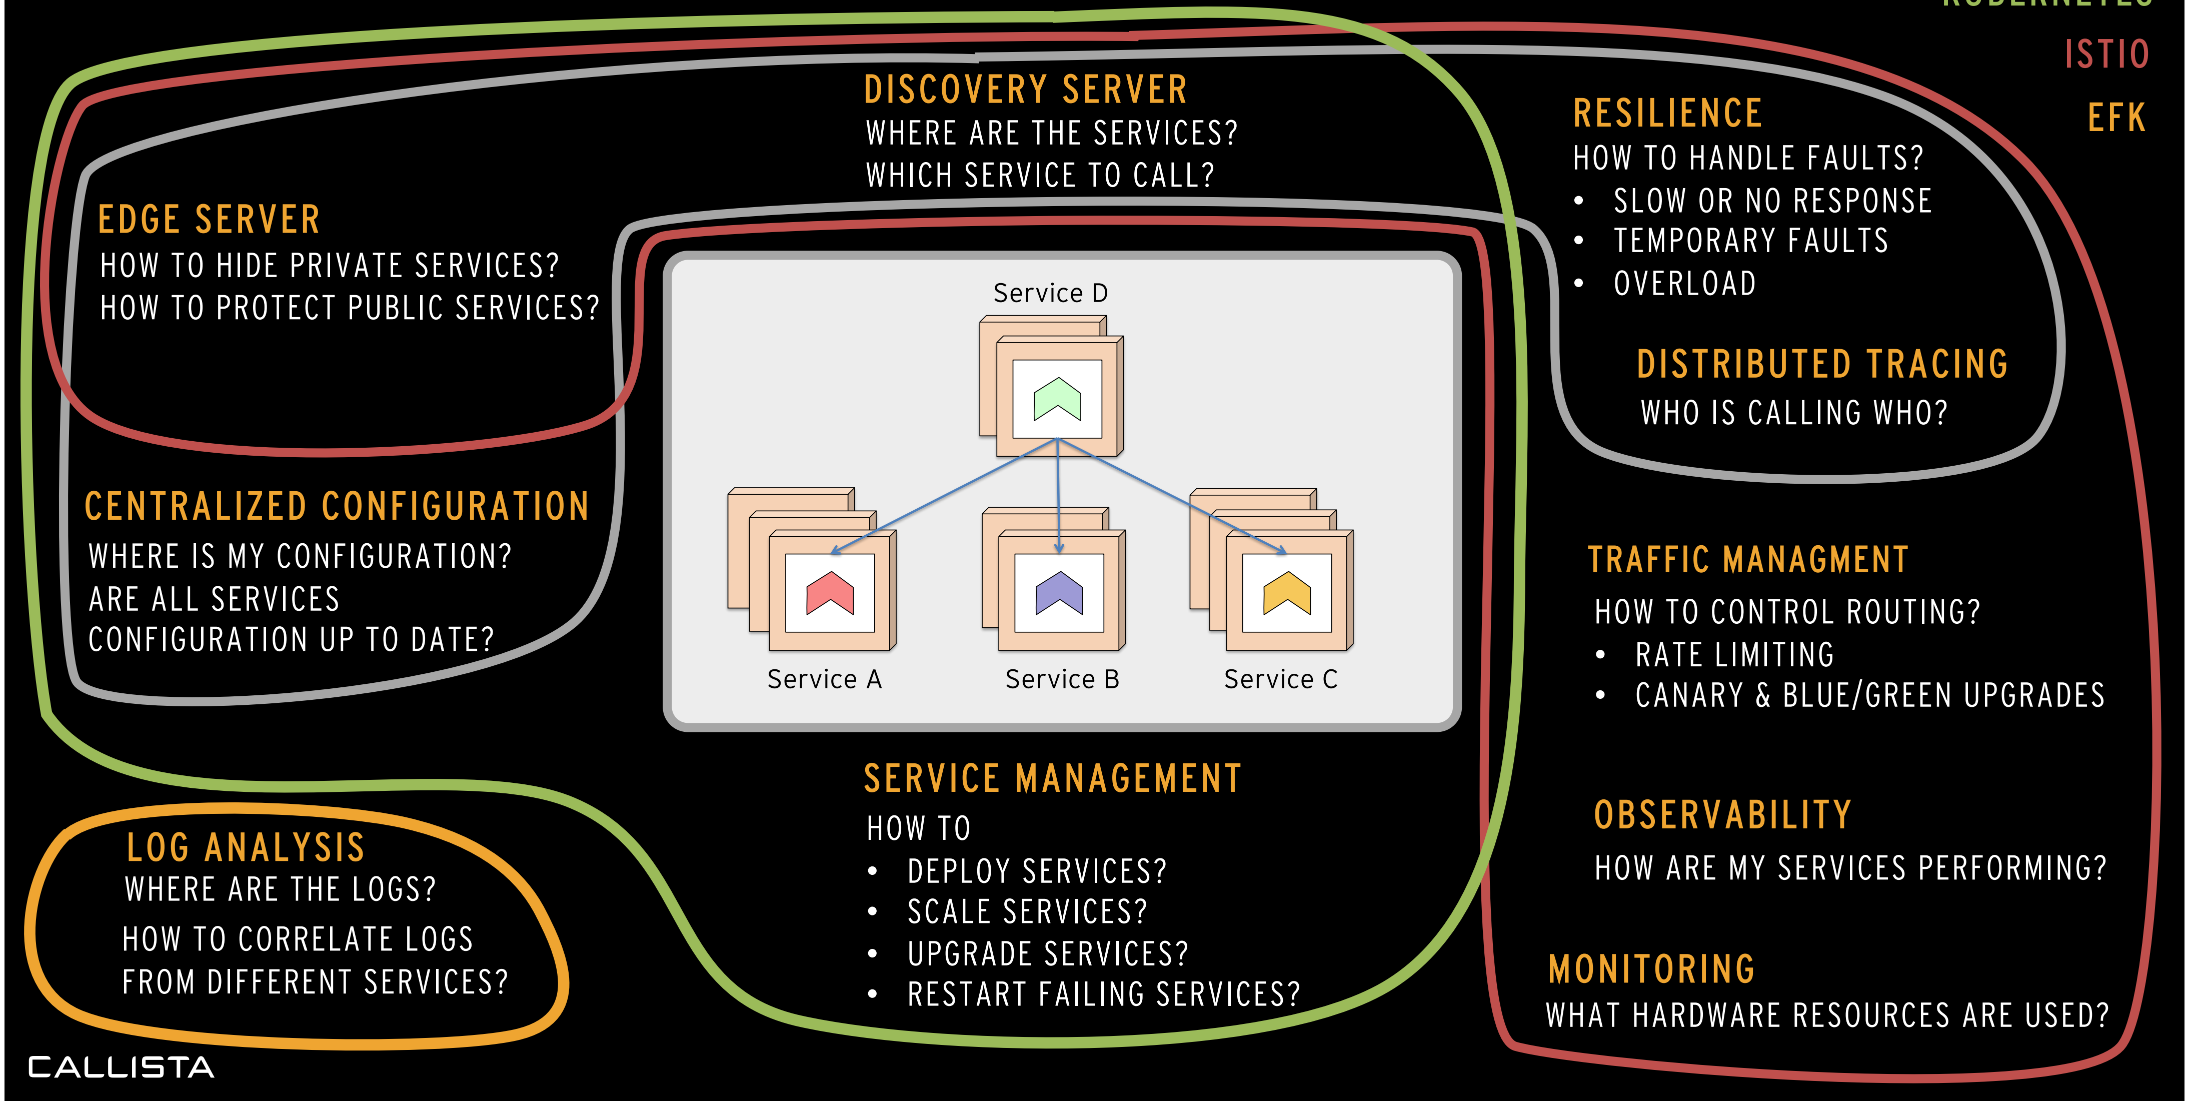

DISCOVERY SERVER

WHERE ARE THE SERVICES?
WHICH SERVICE TO CALL?

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?
HOW TO PROTECT PUBLIC SERVICES?

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?
ARE ALL SERVICES CONFIGURATION UP TO DATE?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS FROM DIFFERENT SERVICES?

SERVICE MANAGEMENT

HOW TO

- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

RESILIENCE

HOW TO HANDLE FAULTS?

- SLOW OR NO RESPONSE
- TEMPORARY FAULTS
- OVERLOAD

DISTRIBUTED TRACING

WHO IS CALLING WHO?

TRAFFIC MANAGMENT

HOW TO CONTROL ROUTING?

- RATE LIMITING
- CANARY & BLUE/GREEN UPGRADES

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CALLISTA

WHERE ARE WE?

- Why?
- Challenges
- Open Source to the rescue!
- Overlaps
- Demo
- Summary

OVERLAPS

Capability	Spring Cloud	Kubernetes	Istio
Service Discovery	X	X	
Central Configuration	X	X	
Edge Server	X	X	X
Distributed Tracing	X		X
Resilience	X		X

FEATURE COMPLETENESS, E.G. FOR AN EDGE SERVER

Feature	Spring Cloud Gateway	Kubernetes Ingress Controller	Istio Ingress Gateway
Security			
- OAuth 2.0 & OIDC	X	X	X
- Automated provisioning and renewal of certificates		X	X
Routing			
- URL path based	X	X	X
- Header based	X		X
Observability			X
Traffic Management			X

OVERLAPS - HOW TO CHOOSE?

- Prefer platform over application library
  - Independence of microservice implementations
    - » E.g. language or frameworks

- Exceptions, i.e. use application library for

1. Managing trace ids in a microservice
   - » Setting inbound trace id on outbound requests
2. Resilience mechanisms, e.g. timeout, retry and circuit breakers
   - » Fine tuning often depends on business logic

   Note: Platform based resilience is much better than none at al…

OVERLAPS - SELECTIONS

Capability	Spring Cloud	Kubernetes	Istio
Service Discovery	X	**X**	
Central Configuration	X	**X**	
Edge Server	X	X	**X**
Distributed Tracing	**X**		**X**
Resilience	**X**		X

OVERLAPS - SELECTIONS

Capability	Spring Cloud	Kubernetes	Istio
Service Discovery	Netflix Eureka Spring Cloud Load Balancer	Kube Proxy & Service objects	
Central Configuration	Spring Cloud Config server	Config Maps & Secrets	
Edge Server	Spring Cloud Gateway	Ingress Controller	Ingress Gateway
Distributed Tracing	• Spring Cloud Sleuth • Zipkin		• Jaeger • Zipkin
Resilience	Resilience4J		Timeout, Retries & Outlier Detection

CAPABILITY MAPPING

SPRING CLOUD

KUBERNETES

ISTIO

EFK

EDGE SERVER

HOW TO HIDE PRIVATE SERVICES?
HOW TO PROTECT PUBLIC SERVICES?

DISCOVERY SERVER

WHERE ARE THE SERVICES?
WHICH SERVICE TO CALL?

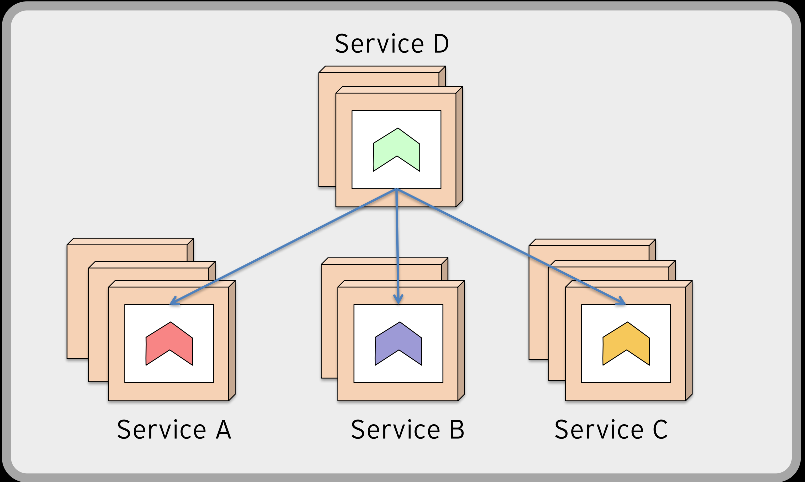

CENTRALIZED CONFIGURATION

WHERE IS MY CONFIGURATION?
ARE ALL SERVICES
CONFIGURATION UP TO DATE?

SERVICE MANAGEMENT

HOW TO
- DEPLOY SERVICES?
- SCALE SERVICES?
- UPGRADE SERVICES?
- RESTART FAILING SERVICES?

LOG ANALYSIS

WHERE ARE THE LOGS?

HOW TO CORRELATE LOGS
FROM DIFFERENT SERVICES?

RESILIENCE

HOW TO HANDLE FAULTS?
- SLOW OR NO RESPONSE
- TEMPORARY FAULTS
- OVERLOAD

DISTRIBUTED TRACING

WHO IS CALLING WHO?

TRAFFIC MANAGMENT

HOW TO CONTROL ROUTING?
- RATE LIMITING
- CANARY & BLUE/GREEN UPGRADES

OBSERVABILITY

HOW ARE MY SERVICES PERFORMING?

MONITORING

WHAT HARDWARE RESOURCES ARE USED?

CALLISTA

SPRING CLOUD + KUBERNETES + ISTIO

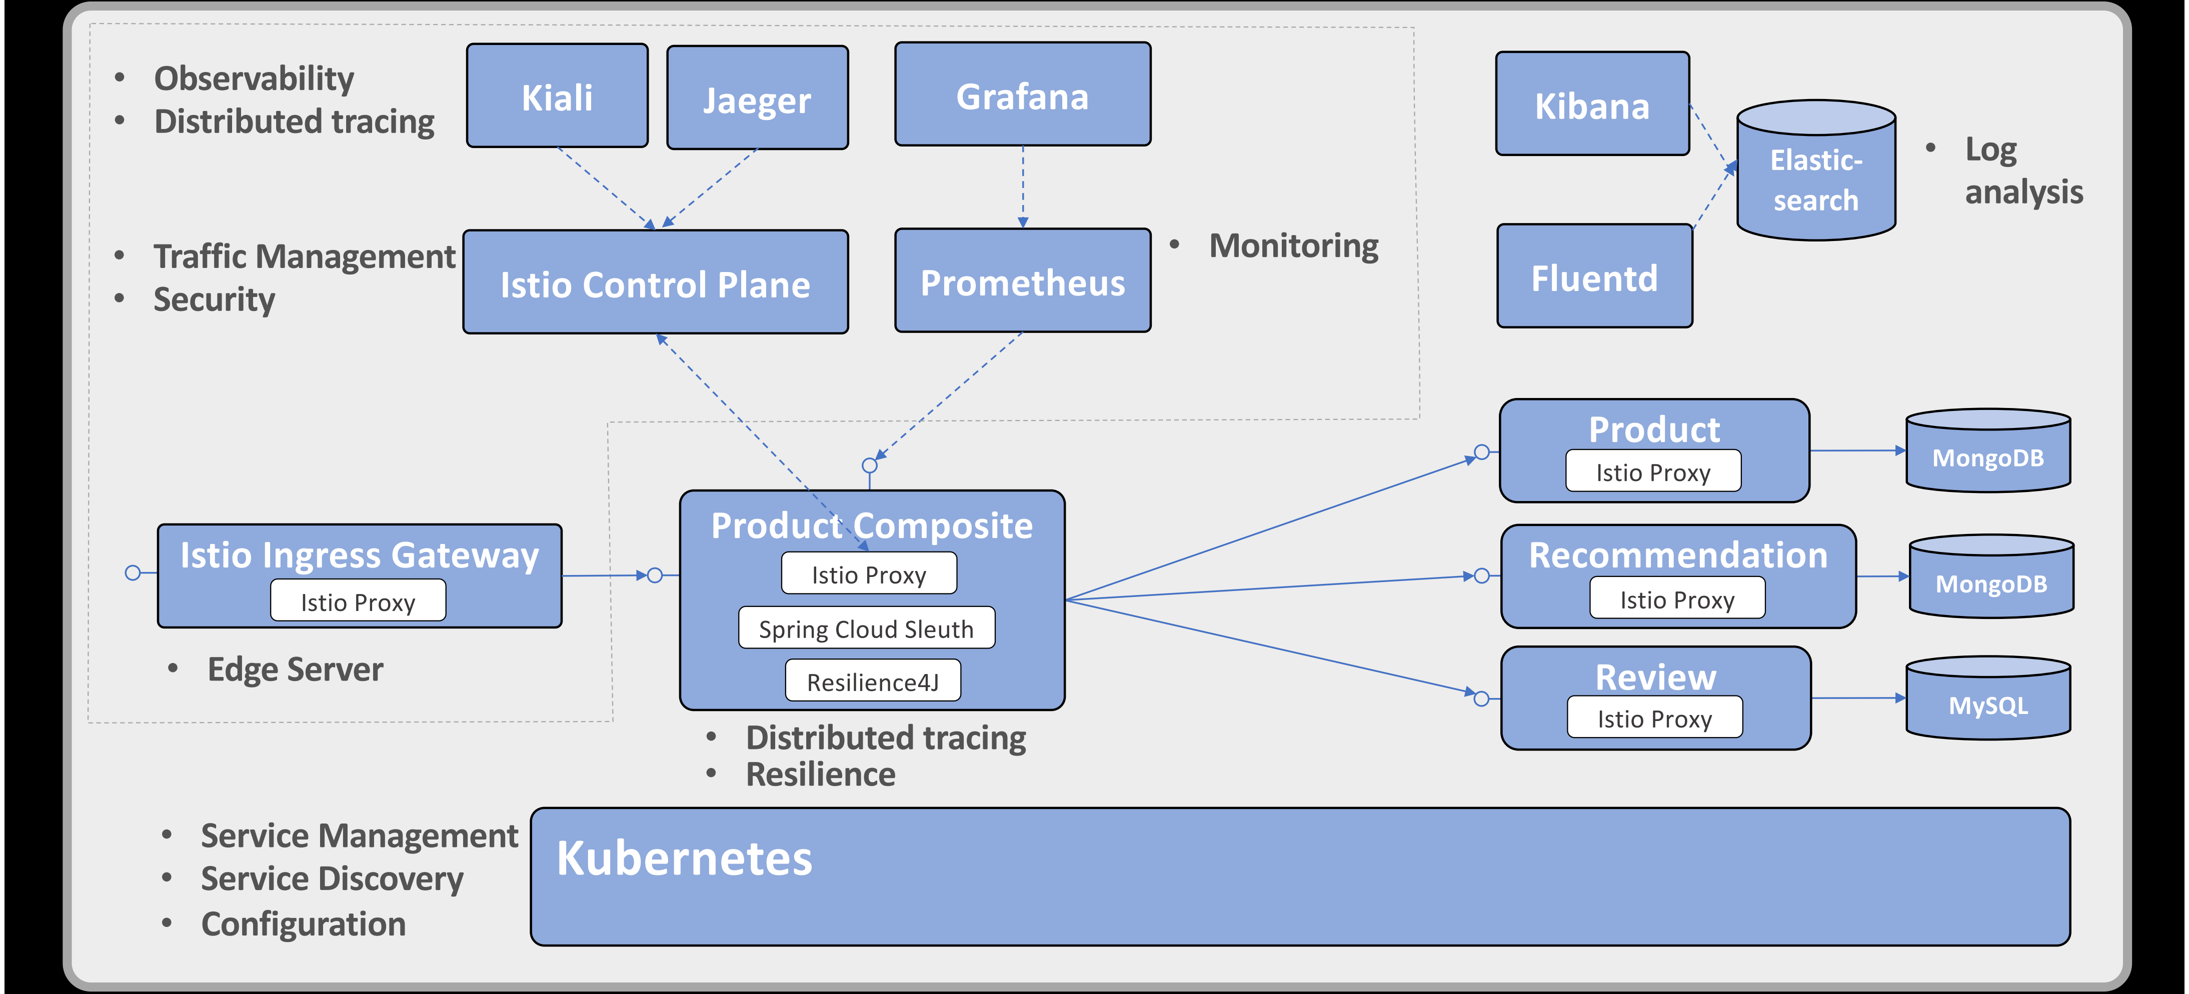

WHERE ARE WE?

- Why?
- Challenges
- Open Source to the rescue!
- Overlaps
- **Demo**
  - Observability
  - Logging
  - Tracing
  - Monitoring
  - Resilience
- Summary

DEMO - OBSERVABILITY

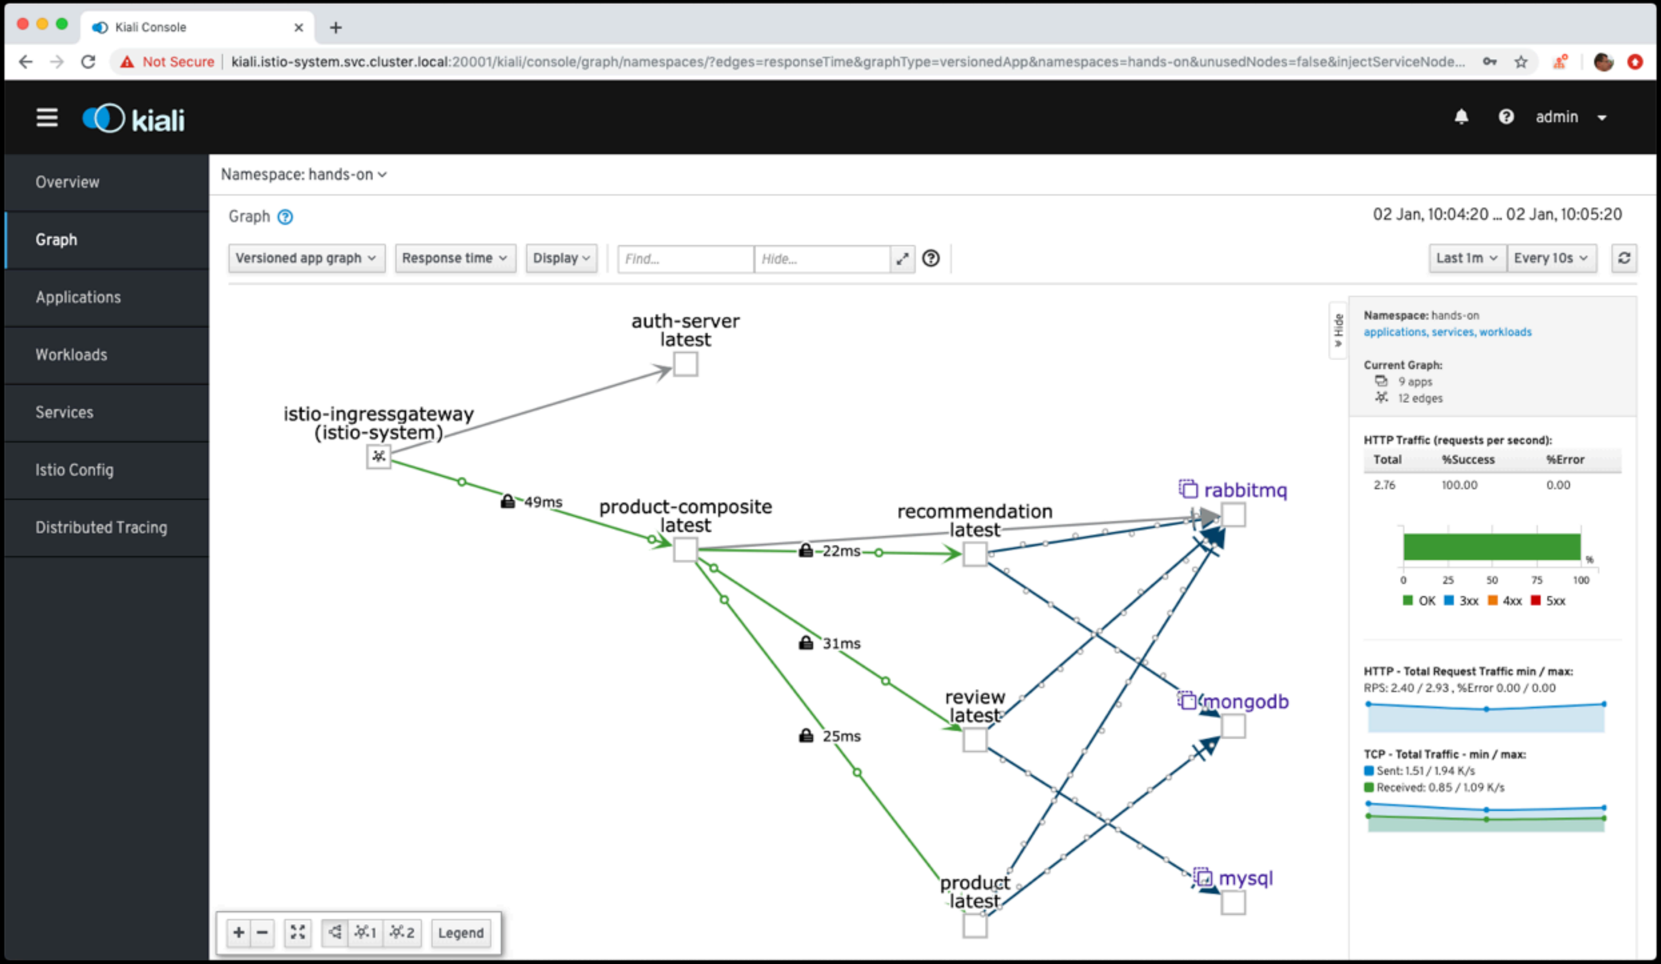

Kiali Console

Not Secure | kiali.istio-system.svc.cluster.local:20001/kiali/console/graph/namespaces/?edges=responseTime&graphType=versionedApp&namespaces=hands-on&unusedNodes=false&injectServiceNode...

Namespace: hands-on

Graph

02 Jan, 10:04:20 ... 02 Jan, 10:05:20

Versioned app graph | Response time | Display | Find... | Hide... | Last 1m | Every 10s

Namespace: hands-on
applications, services, workloads

Current Graph:
9 apps
12 edges

HTTP Traffic (requests per second):

Total	%Success	%Error
2.76	100.00	0.00

HTTP - Total Request Traffic min / max:
RPS: 2.40 / 2.93 , %Error 0.00 / 0.00

TCP - Total Traffic - min / max:
Sent: 1.51 / 1.94 K/s
Received: 0.85 / 1.09 K/s

DEMO - CENTRALIZED LOGGING

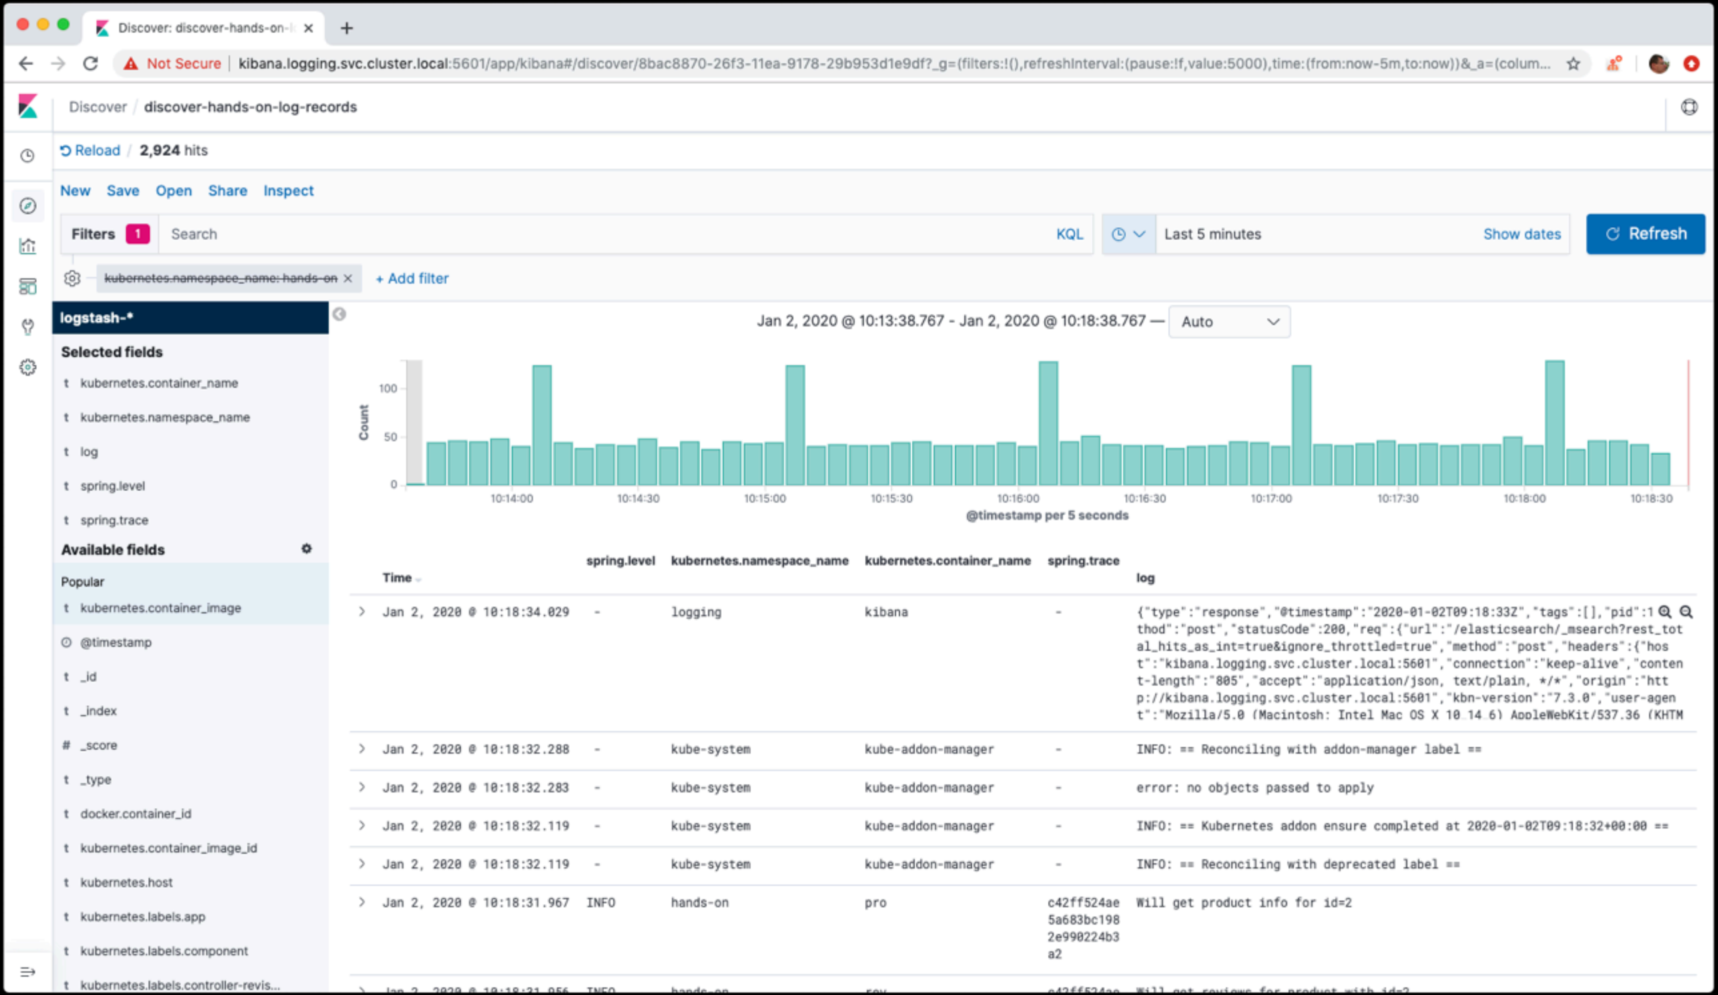

kibana.logging.svc.cluster.local:5601/app/kibana#/discover/8bac8870-26f3-11ea-9178-29b953d1e9df?_g=(filters:!(),refreshInterval:(pause:!f,value:5000),time:(from:now-5m,to:now))&_a=(colum...

Discover / discover-hands-on-log-records

Reload / 2,924 hits

Last 5 minutes

~~kubernetes.namespace_name: hands-on~~

Jan 2, 2020 @ 10:13:38.767 - Jan 2, 2020 @ 10:18:38.767 —

Time	spring.level	kubernetes.namespace_name	kubernetes.container_name	spring.trace	log
Jan 2, 2020 @ 10:18:34.029	-	logging	kibana	-	{"type":"response","@timestamp":"2020-01-02T09:18:33Z","tags":[],"pid":1 thod":"post","statusCode":200,"req":{"url":"/elasticsearch/_msearch?rest_total_hits_as_int=true&ignore_throttled=true","method":"post","headers":{"host":"kibana.logging.svc.cluster.local:5601","connection":"keep-alive","content-length":"805","accept":"application/json, text/plain, */*","origin":"http://kibana.logging.svc.cluster.local:5601","kbn-version":"7.3.0","user-agent":"Mozilla/5.0 (Macintosh; Intel Mac OS X 10_14_6) AppleWebKit/537.36 (KHTM
Jan 2, 2020 @ 10:18:32.288	-	kube-system	kube-addon-manager	-	INFO: == Reconciling with addon-manager label ==
Jan 2, 2020 @ 10:18:32.283	-	kube-system	kube-addon-manager	-	error: no objects passed to apply
Jan 2, 2020 @ 10:18:32.119	-	kube-system	kube-addon-manager	-	INFO: == Kubernetes addon ensure completed at 2020-01-02T09:18:32+00:00 ==
Jan 2, 2020 @ 10:18:32.119	-	kube-system	kube-addon-manager	-	INFO: == Reconciling with deprecated label ==
Jan 2, 2020 @ 10:18:31.967	INFO	hands-on	pro	c42ff524ae5a683bc1982e990224b3a2	Will get product info for id=2
Jan 2, 2020 @ 10:18:31.956	INFO	hands-on	rev	c42ff524ae	Will get reviews for product with id=2

CALLISTA

DEMO - CENTRALIZED LOGGING

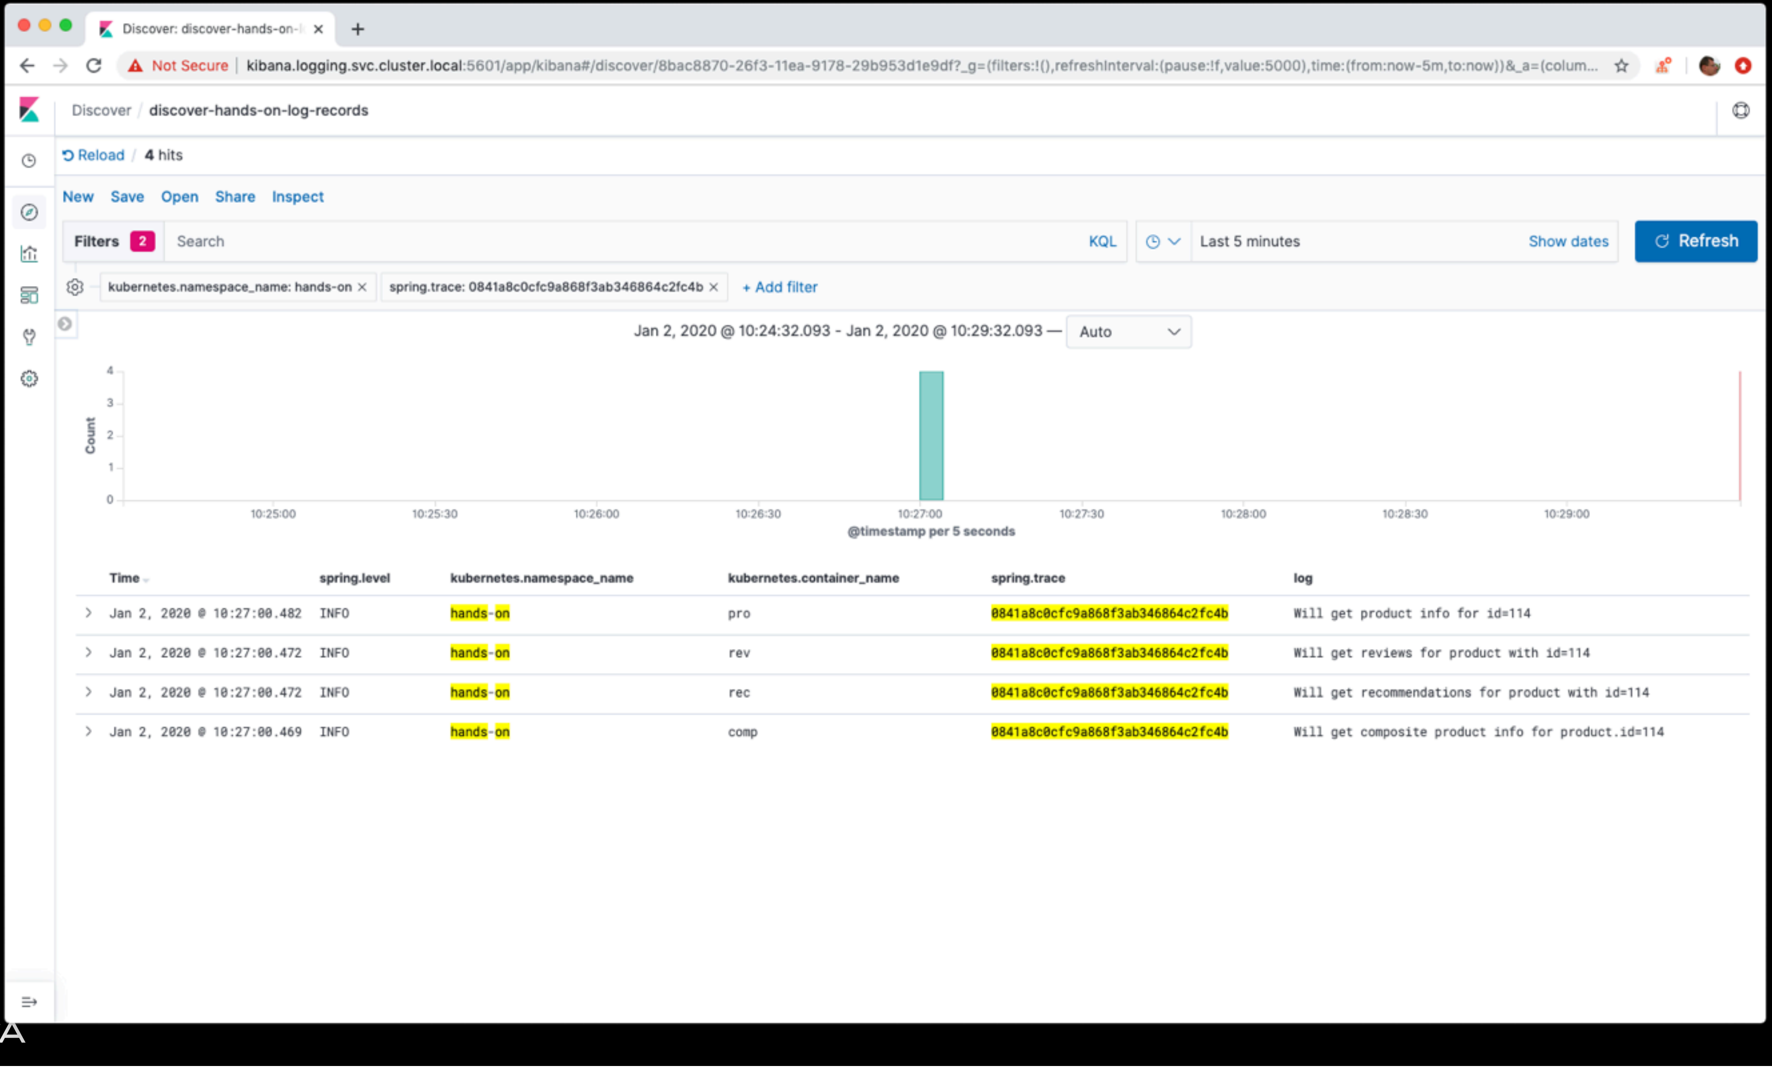

kibana.logging.svc.cluster.local:5601/app/kibana#/discover/8bac8870-26f3-11ea-9178-29b953d1e9df?_g=(filters:!(),refreshInterval:(pause:!f,value:5000),time:(from:now-5m,to:now))&_a=(colum...

Last 5 minutes

kubernetes.namespace_name: hands-on ×

spring.trace: 0841a8c0cfc9a868f3ab346864c2fc4b ×

Jan 2, 2020 @ 10:24:32.093 - Jan 2, 2020 @ 10:29:32.093 —

Time	spring.level	kubernetes.namespace_name	kubernetes.container_name	spring.trace	log
Jan 2, 2020 @ 10:27:00.482	INFO	hands-on	pro	0841a8c0cfc9a868f3ab346864c2fc4b	Will get product info for id=114
Jan 2, 2020 @ 10:27:00.472	INFO	hands-on	rev	0841a8c0cfc9a868f3ab346864c2fc4b	Will get reviews for product with id=114
Jan 2, 2020 @ 10:27:00.472	INFO	hands-on	rec	0841a8c0cfc9a868f3ab346864c2fc4b	Will get recommendations for product with id=114
Jan 2, 2020 @ 10:27:00.469	INFO	hands-on	comp	0841a8c0cfc9a868f3ab346864c2fc4b	Will get composite product info for product.id=114

DEMO - DISTRIBUTED TRACING

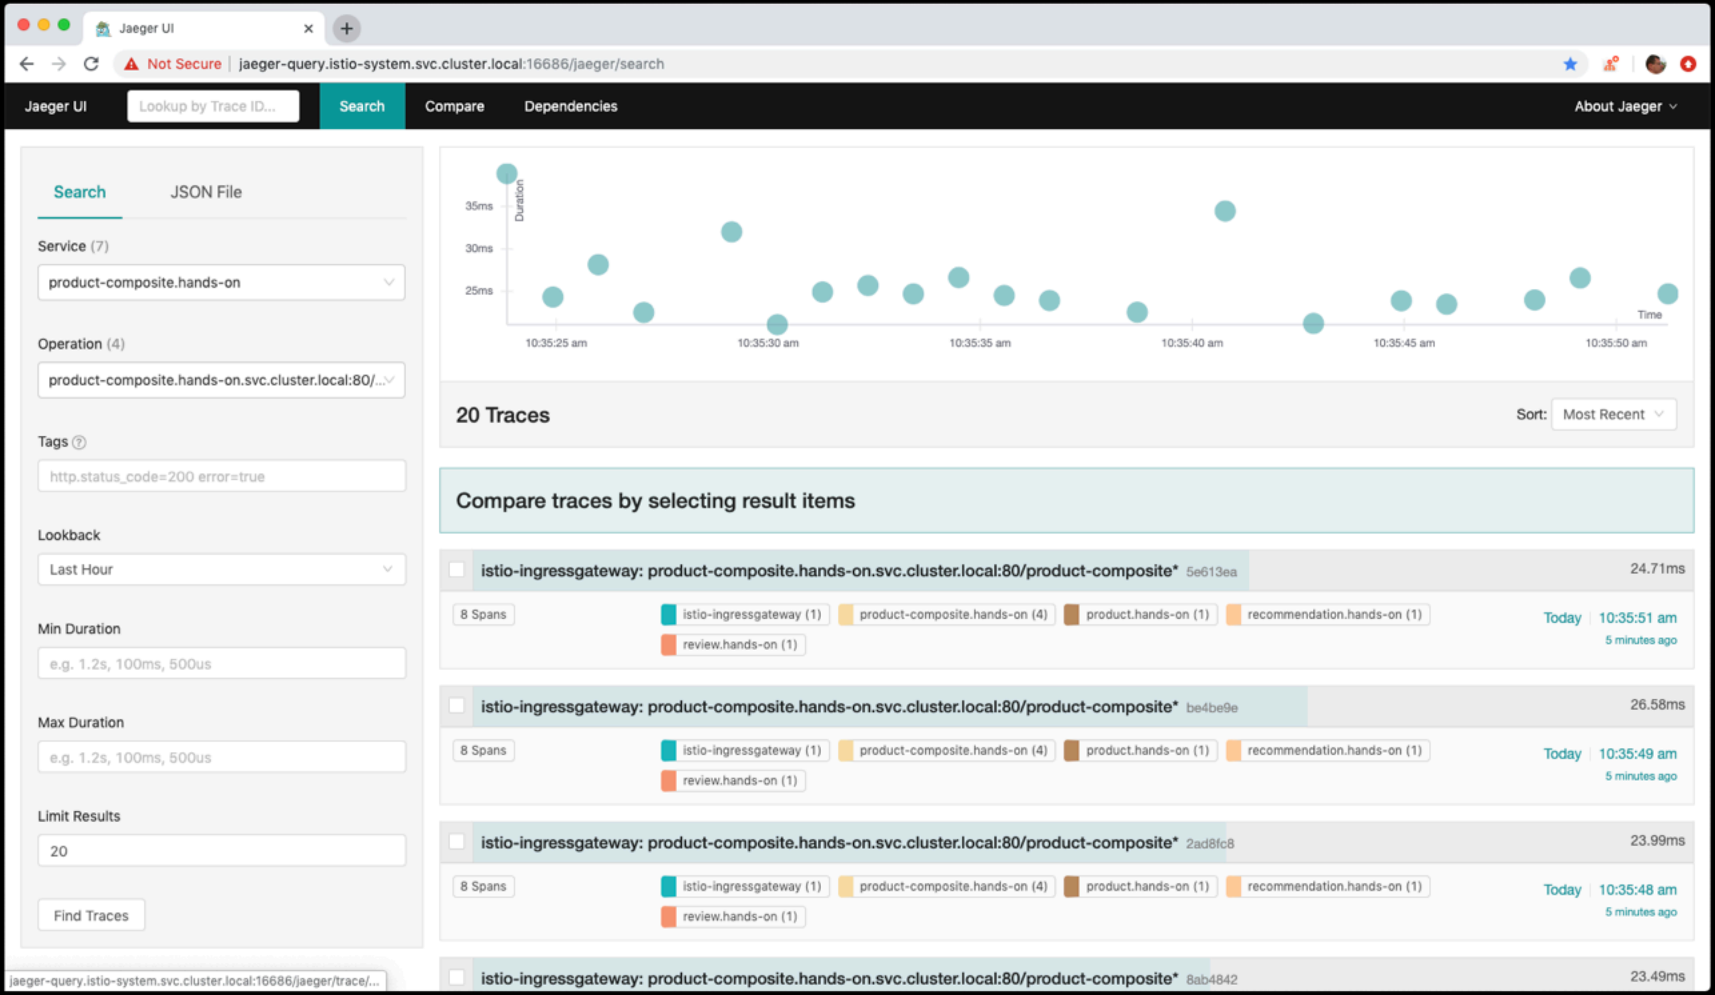

Search | JSON File

Service (7)
product-composite.hands-on

Operation (4)
product-composite.hands-on.svc.cluster.local:80/...

Tags
http.status_code=200 error=true

Lookback
Last Hour

Min Duration
e.g. 1.2s, 100ms, 500us

Max Duration
e.g. 1.2s, 100ms, 500us

Limit Results
20

Find Traces

20 Traces — Sort: Most Recent

Compare traces by selecting result items

□ istio-ingressgateway: product-composite.hands-on.svc.cluster.local:80/product-composite* 5e613ea — 24.71ms
8 Spans — istio-ingressgateway (1), product-composite.hands-on (4), product.hands-on (1), recommendation.hands-on (1), review.hands-on (1) — Today | 10:35:51 am, 5 minutes ago

□ istio-ingressgateway: product-composite.hands-on.svc.cluster.local:80/product-composite* be4be9e — 26.58ms
8 Spans — istio-ingressgateway (1), product-composite.hands-on (4), product.hands-on (1), recommendation.hands-on (1), review.hands-on (1) — Today | 10:35:49 am, 5 minutes ago

□ istio-ingressgateway: product-composite.hands-on.svc.cluster.local:80/product-composite* 2ad8fc8 — 23.99ms
8 Spans — istio-ingressgateway (1), product-composite.hands-on (4), product.hands-on (1), recommendation.hands-on (1), review.hands-on (1) — Today | 10:35:48 am, 5 minutes ago

□ istio-ingressgateway: product-composite.hands-on.svc.cluster.local:80/product-composite* 8ab4842 — 23.49ms

DEMO - DISTRIBUTED TRACING

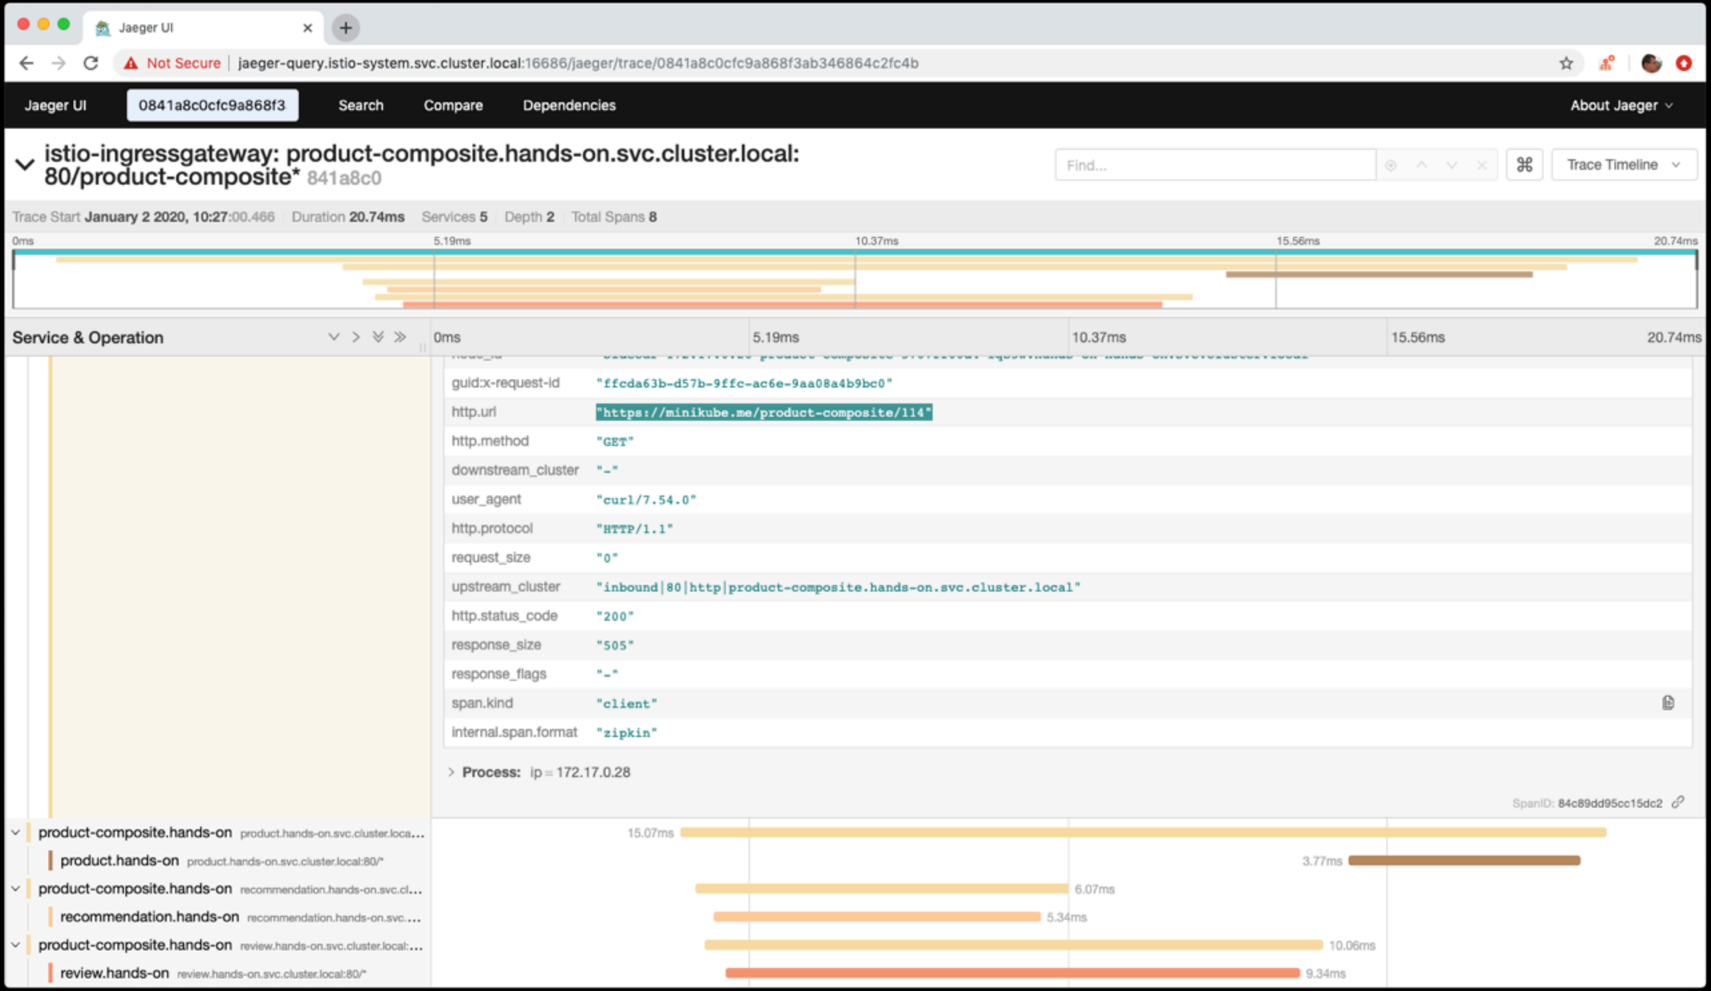

istio-ingressgateway: product-composite.hands-on.svc.cluster.local:80/product-composite* 841a8c0

Trace Start **January 2 2020, 10:27**:00.466 | Duration **20.74ms** | Services **5** | Depth **2** | Total Spans **8**

guid:x-request-id	"ffcda63b-d57b-9ffc-ac6e-9aa08a4b9bc0"
http.url	"https://minikube.me/product-composite/114"
http.method	"GET"
downstream_cluster	"-"
user_agent	"curl/7.54.0"
http.protocol	"HTTP/1.1"
request_size	"0"
upstream_cluster	"inbound|80|http|product-composite.hands-on.svc.cluster.local"
http.status_code	"200"
response_size	"505"
response_flags	"-"
span.kind	"client"
internal.span.format	"zipkin"

› **Process:** ip = 172.17.0.28

SpanID: 84c89dd95cc15dc2

product-composite.hands-on product.hands-on.svc.cluster.loca...	15.07ms
product.hands-on product.hands-on.svc.cluster.local:80/*	3.77ms
product-composite.hands-on recommendation.hands-on.svc.cl...	6.07ms
recommendation.hands-on recommendation.hands-on.svc....	5.34ms
product-composite.hands-on review.hands-on.svc.cluster.local:...	10.06ms
review.hands-on review.hands-on.svc.cluster.local:80/*	9.34ms

CALLISTA

DEMO - MONITORING

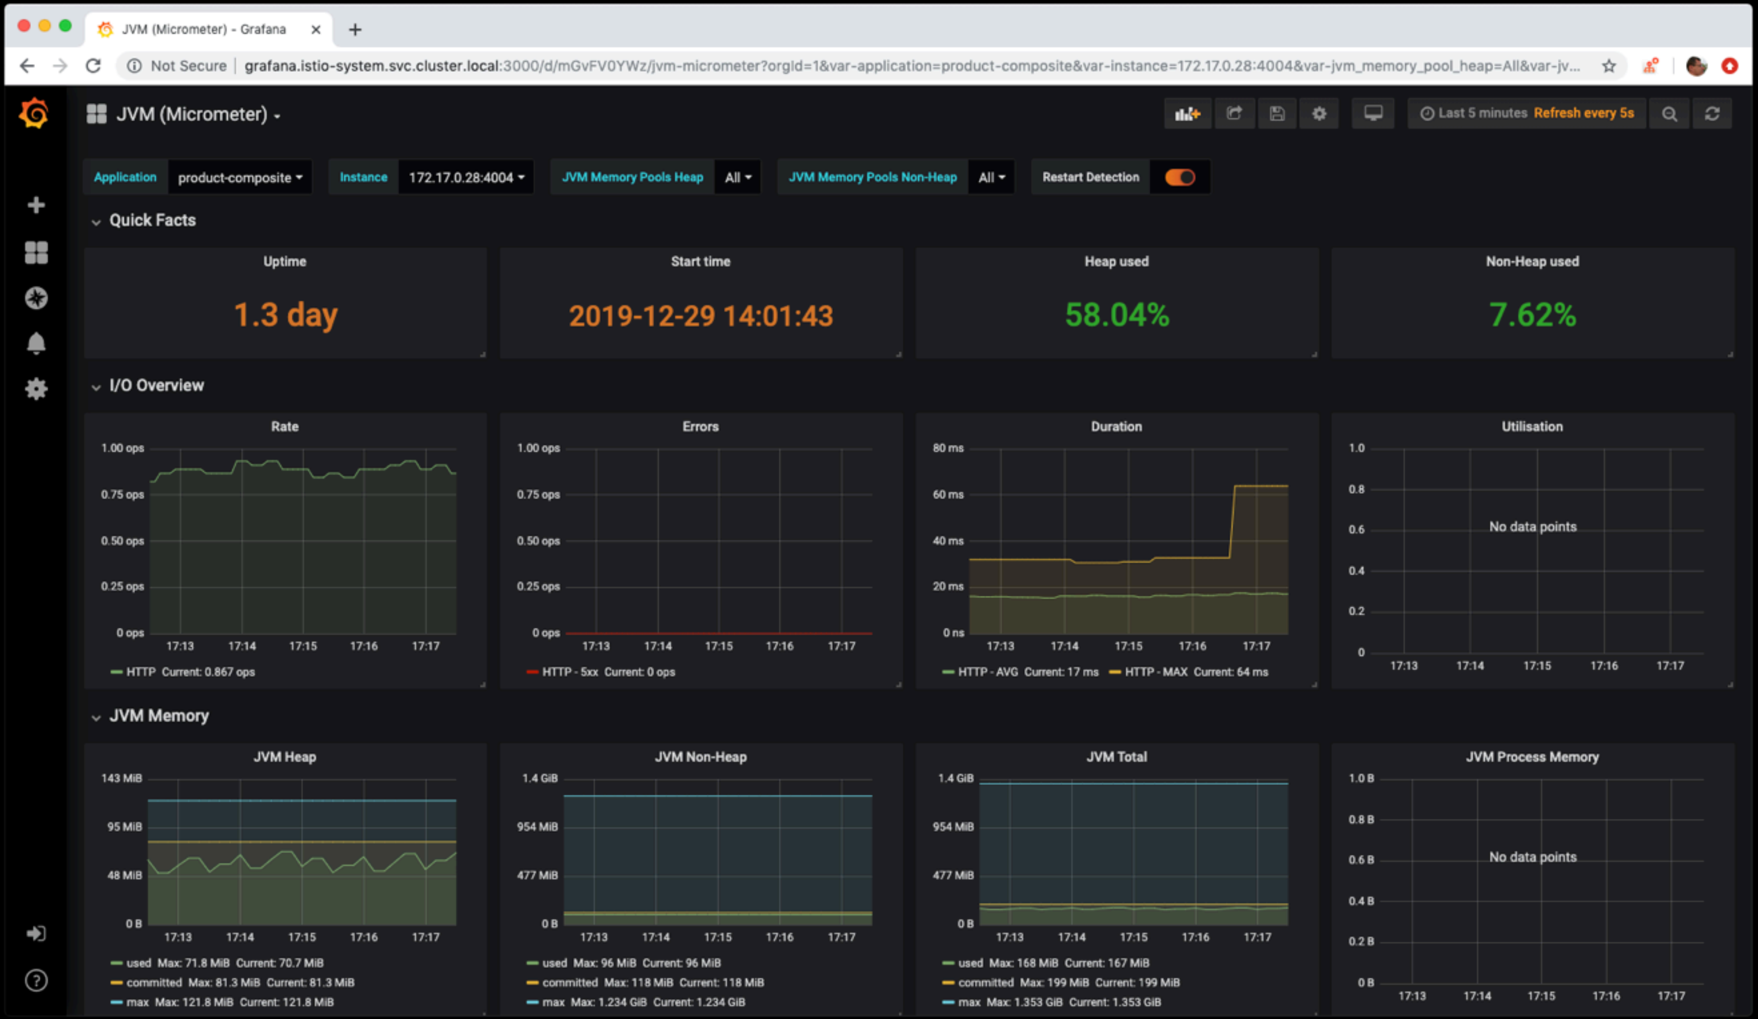

JVM (Micrometer) - Grafana

Not Secure | grafana.istio-system.svc.cluster.local:3000/d/mGvFV0YWz/jvm-micrometer?orgId=1&var-application=product-composite&var-instance=172.17.0.28:4004&var-jvm_memory_pool_heap=All&var-jv...

JVM (Micrometer)

Last 5 minutes Refresh every 5s

Application: product-composite | Instance: 172.17.0.28:4004 | JVM Memory Pools Heap: All | JVM Memory Pools Non-Heap: All | Restart Detection

Quick Facts

Uptime: 1.3 day

Start time: 2019-12-29 14:01:43

Heap used: 58.04%

Non-Heap used: 7.62%

I/O Overview

Rate — HTTP Current: 0.867 ops

Errors — HTTP - 5xx Current: 0 ops

Duration — HTTP - AVG Current: 17 ms — HTTP - MAX Current: 64 ms

Utilisation — No data points

JVM Memory

JVM Heap
- used Max: 71.8 MiB Current: 70.7 MiB
- committed Max: 81.3 MiB Current: 81.3 MiB
- max Max: 121.8 MiB Current: 121.8 MiB

JVM Non-Heap
- used Max: 96 MiB Current: 96 MiB
- committed Max: 118 MiB Current: 118 MiB
- max Max: 1.234 GiB Current: 1.234 GiB

JVM Total
- used Max: 168 MiB Current: 167 MiB
- committed Max: 199 MiB Current: 199 MiB
- max Max: 1.353 GiB Current: 1.353 GiB

JVM Process Memory — No data points

DEMO - RESILIENCE

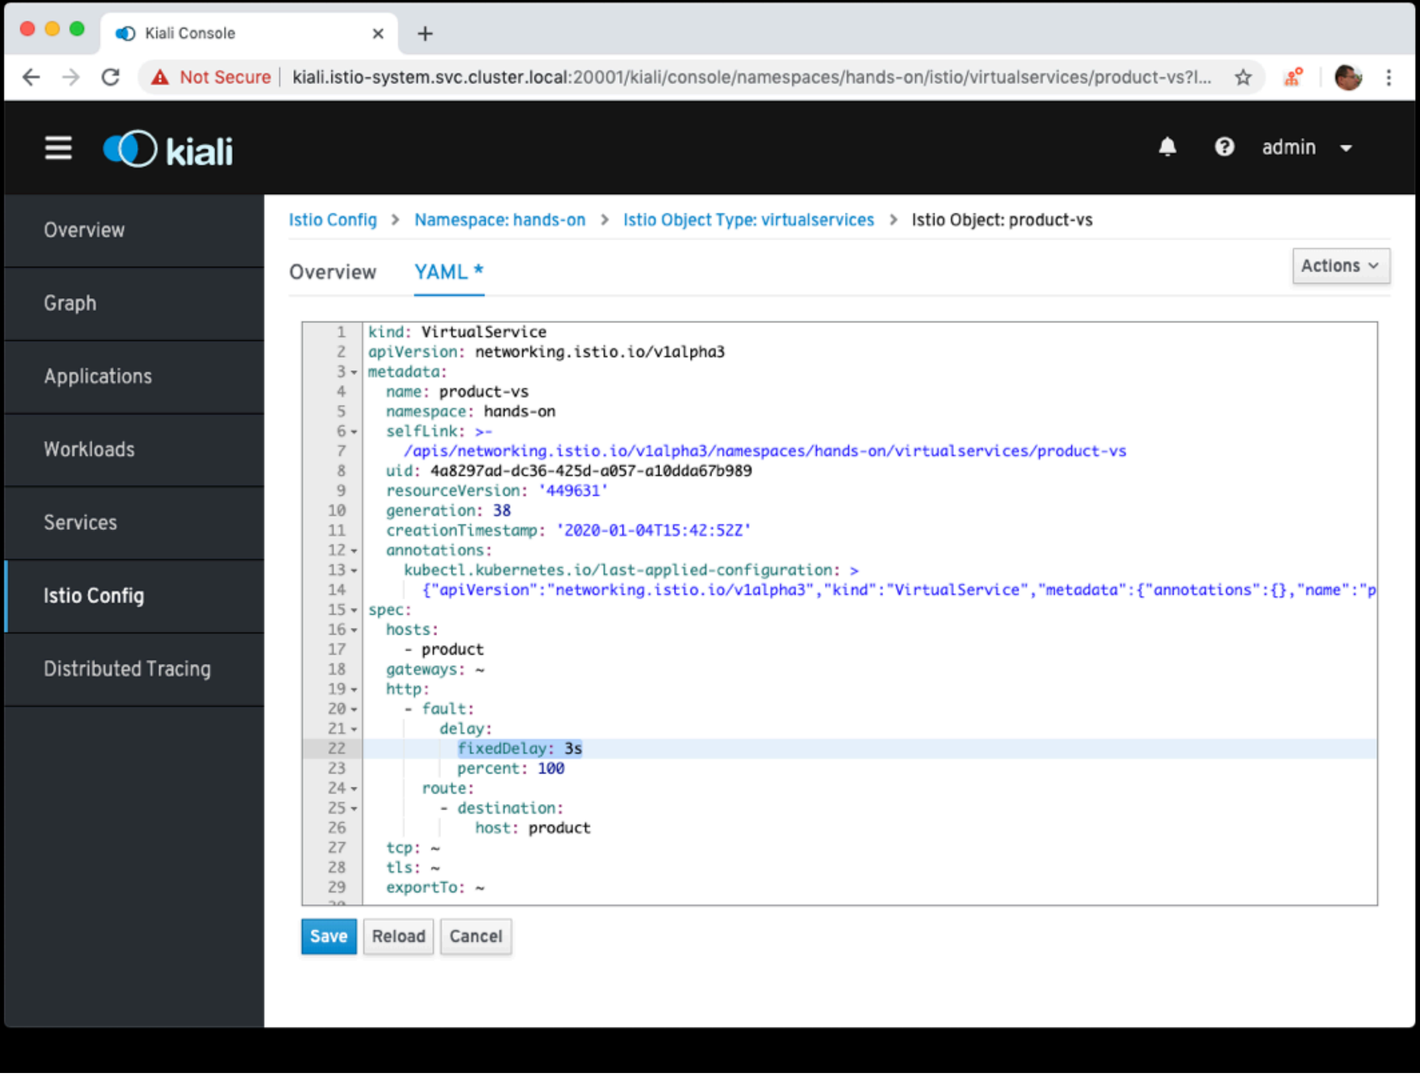
Kiali Console
Not Secure
kiali.istio-system.svc.cluster.local:20001/kiali/console/namespaces/hands-on/istio/virtualservices/product-vs?l...
kiali
admin
Overview
Graph
Applications
Workloads
Services
Istio Config
Distributed Tracing
Istio Config > Namespace: hands-on > Istio Object Type: virtualservices > Istio Object: product-vs
Overview
YAML *
Actions
Save
Reload
Cancel

```yaml
kind: VirtualService
apiVersion: networking.istio.io/v1alpha3
metadata:
  name: product-vs
  namespace: hands-on
  selfLink: >-
    /apis/networking.istio.io/v1alpha3/namespaces/hands-on/virtualservices/product-vs
  uid: 4a8297ad-dc36-425d-a057-a10dda67b989
  resourceVersion: '449631'
  generation: 38
  creationTimestamp: '2020-01-04T15:42:52Z'
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: >
      {"apiVersion":"networking.istio.io/v1alpha3","kind":"VirtualService","metadata":{"annotations":{},"name":"p
spec:
  hosts:
    - product
  gateways: ~
  http:
    - fault:
        delay:
          fixedDelay: 3s
          percent: 100
      route:
        - destination:
            host: product
  tcp: ~
  tls: ~
  exportTo: ~
```

CALLISTA

DEMO - RESILIENCE

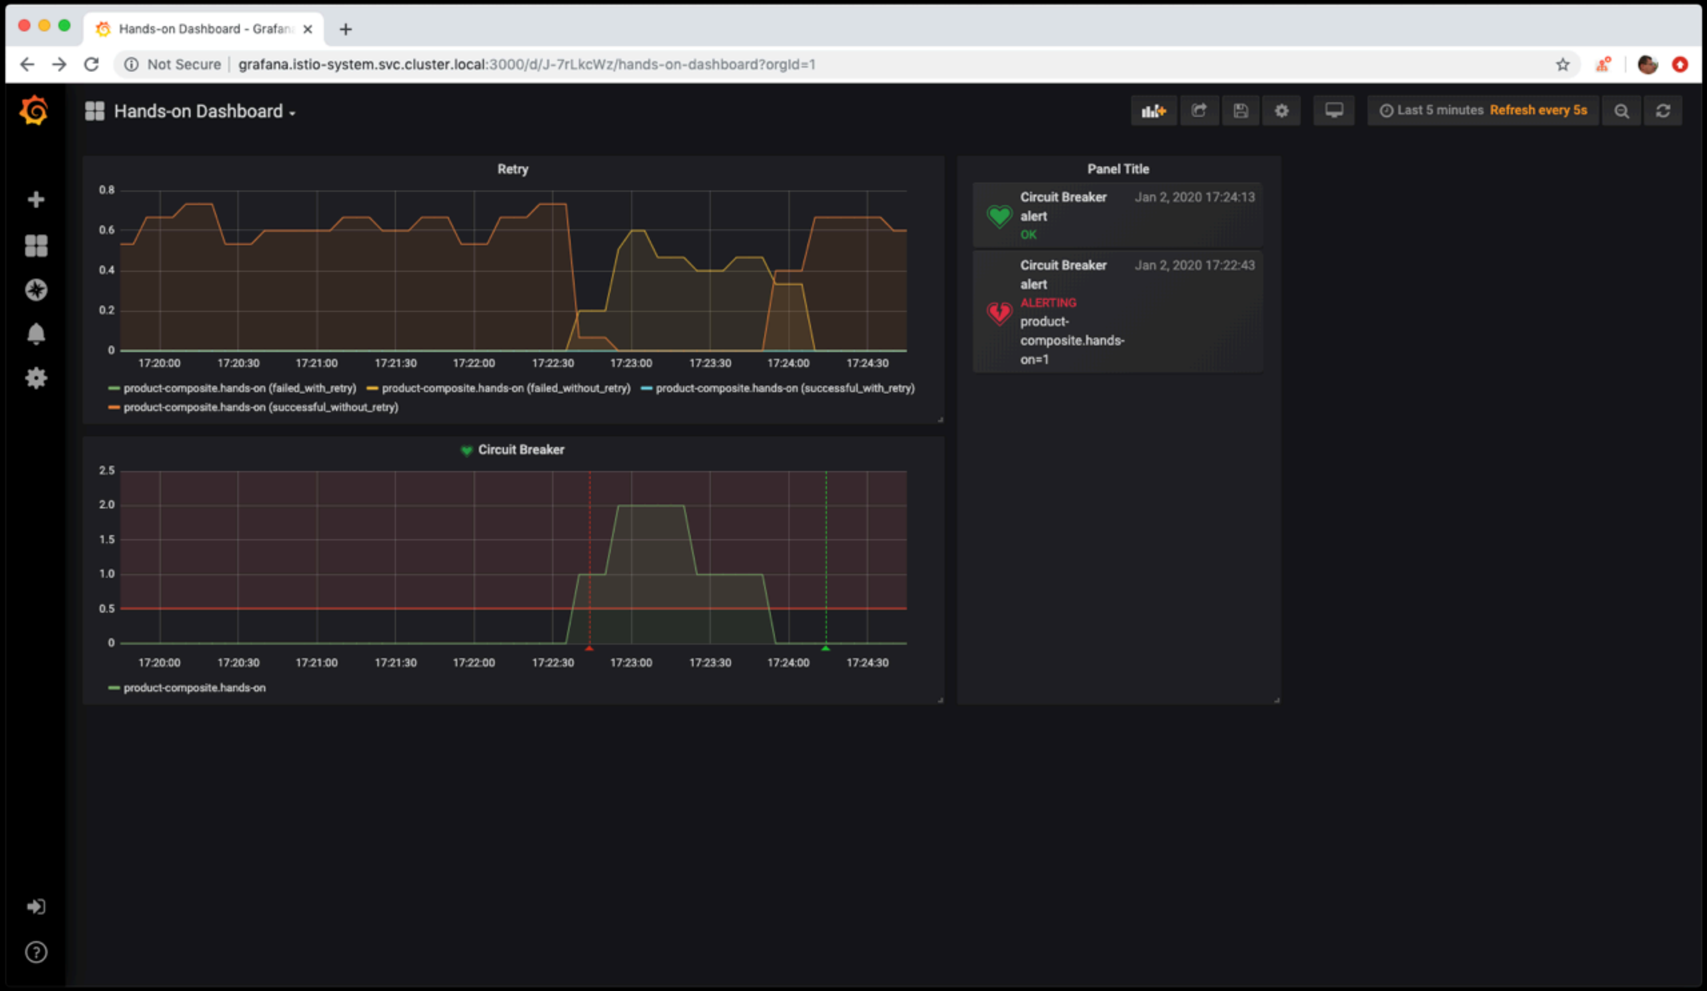
Hands-on Dashboard - Grafan
Not Secure
grafana.istio-system.svc.cluster.local:3000/d/J-7rLkcWz/hands-on-dashboard?orgId=1
Hands-on Dashboard
Last 5 minutes Refresh every 5s
Retry
0.8
0.6
0.4
0.2
0
17:20:00
17:20:30
17:21:00
17:21:30
17:22:00
17:22:30
17:23:00
17:23:30
17:24:00
17:24:30
product-composite.hands-on (failed_with_retry)
product-composite.hands-on (failed_without_retry)
product-composite.hands-on (successful_with_retry)
product-composite.hands-on (successful_without_retry)
Circuit Breaker
2.5
2.0
1.5
1.0
0.5
0
17:20:00
17:20:30
17:21:00
17:21:30
17:22:00
17:22:30
17:23:00
17:23:30
17:24:00
17:24:30
product-composite.hands-on

Panel Title

Circuit Breaker alert Jan 2, 2020 17:24:13
OK

Circuit Breaker alert Jan 2, 2020 17:22:43
ALERTING
product-composite.hands-on=1

DEMO - RESILIENCE

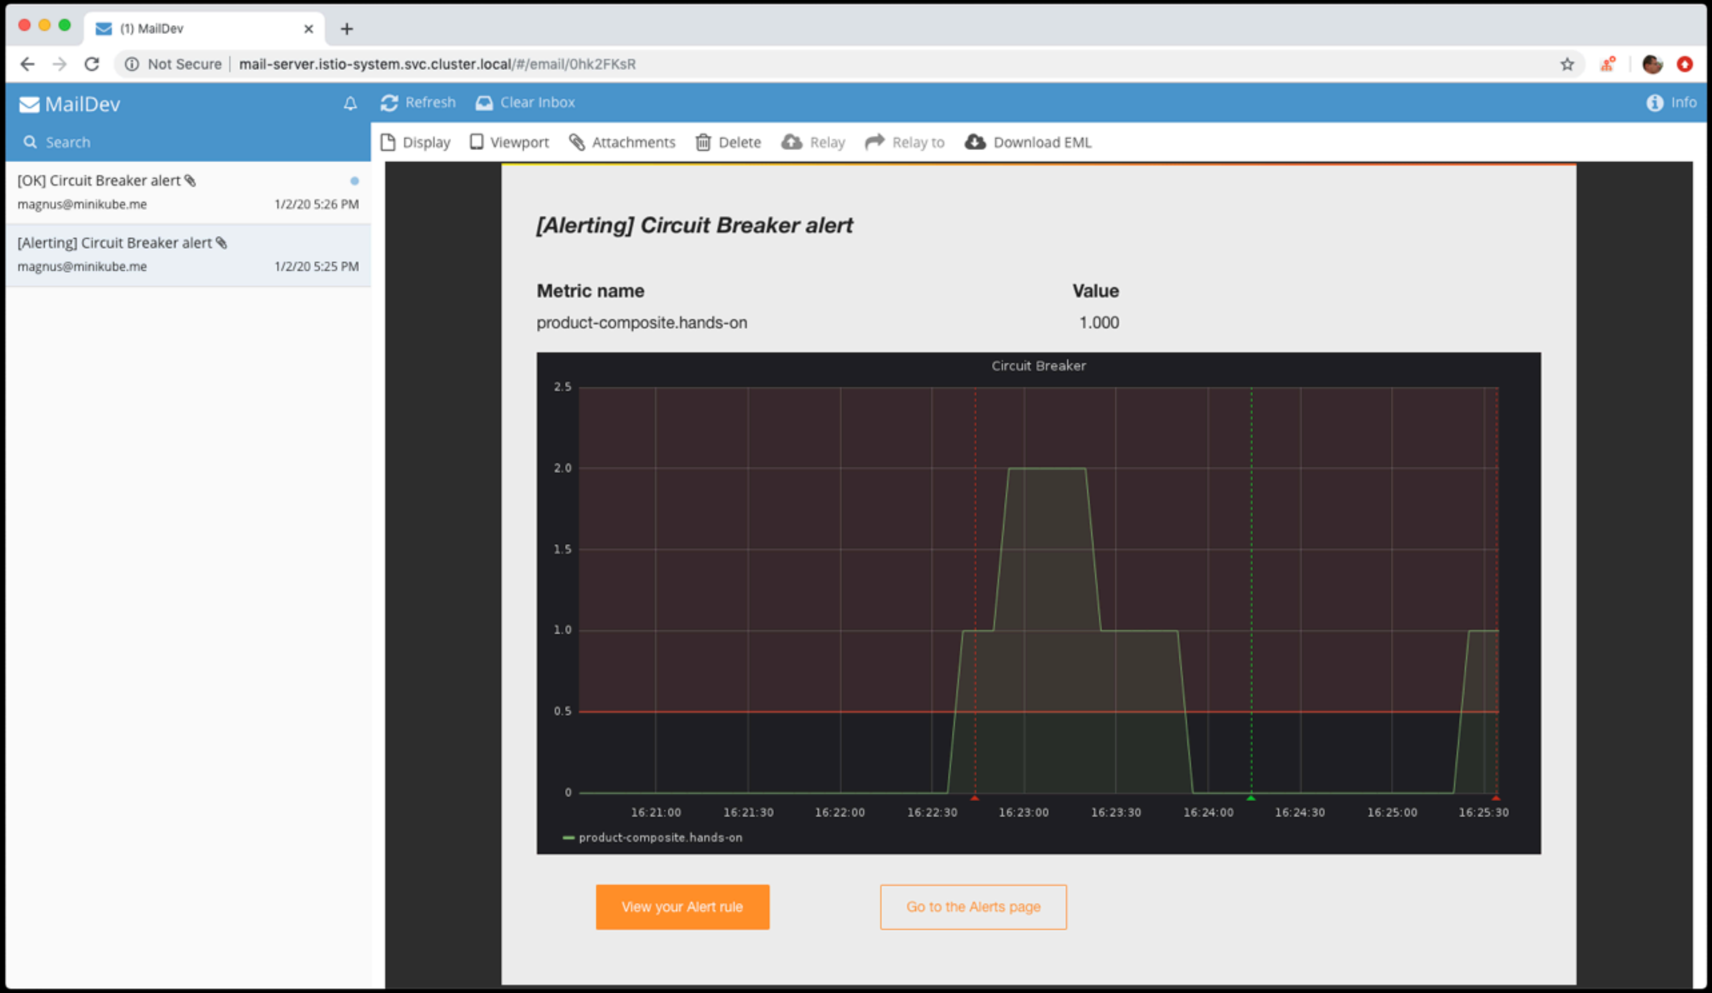

mail-server.istio-system.svc.cluster.local/#/email/0hk2FKsR

[OK] Circuit Breaker alert
magnus@minikube.me 1/2/20 5:26 PM

[Alerting] Circuit Breaker alert
magnus@minikube.me 1/2/20 5:25 PM

[Alerting] Circuit Breaker alert

Metric name	Value
product-composite.hands-on	1.000

SUMMARY

- Microservices promise
  - Easier to scale
  - Faster release cycles
- Cooperating microservices ➔ Distributed System
  - Inherent complexity
  - Can be managed with Open Source
    - » Application library, e.g. Spring Cloud
    - » Container orchestrators, e.g. Kubernetes
    - » Service mesh, e.g. Istio
- Handle overlaps
- Works great together!
  - …if used correctly

RECOMMENDED READING

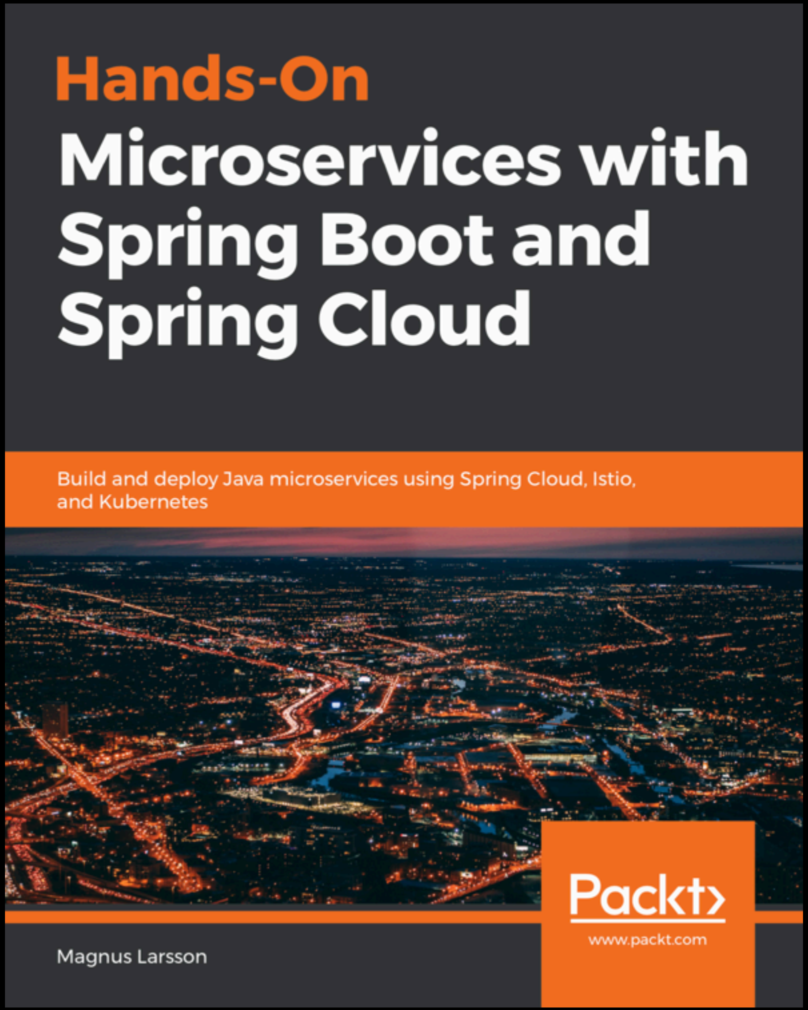
Hands-On
Microservices with Spring Boot and Spring Cloud
Build and deploy Java microservices using Spring Cloud, Istio, and Kubernetes
Packt>
www.packt.com
Magnus Larsson

- Book – Hands-on microservices

https://www.packtpub.com/web-development/hands-on-microservices-with-spring-boot-and-spring-cloud

- Blog series – Java & GO based microservices

https://callistaenterprise.se/blogg/teknik/2015/05/20/blog-series-building-microservices/

CALLISTA