

CONTINUOUS  
**DELIVERY**  
PIPELINE

The logo for 'nexa' is located in the bottom right corner. It consists of a square with a vertical gradient from dark blue on the left to green on the right. The word 'nexa' is written in a white, lowercase, sans-serif font in the center of the square.

*nexa*

# Rodrigo Botti

Backend, Frontend, Ops, Fullstack (?)

**NEXA:** Integrations Team - Crawler, APIs, Ops, Fullstack (?)

**GDG Campinas:** Organizer

*nexa*





# AGENDA

- ▶ Current Scenario
- ▶ Continuous Delivery
- ▶ Canary Release
- ▶ Tooling
- ▶ Future
- ▶ Q&A

**!!! DISCLAIMER !!!**



1.

SCENARIO

Technology  
Platforms

MAIN  
TECHNOLOGIES



**kubernetes**



# MAIN TECHNOLOGIES



**Jenkins**



Spinnaker



Google Cloud Platform

# MICROSERVICES

## Advantages

### MODULARITY

Reduced codebase, easier to reason about, develop, test and deploy.

Single responsibility.

### PARALLELISM

Multiple teams working on separate services.

### SCALABLE

Faster feature delivery.

Independent horizontal/vertical scaling and failure handling.

# MICROSERVICES

## Requirements

### \*\*\* AUTOMATION \*\*\*

Automation is key for having speed, scalability / elasticity specially when dealing with a distributed services application.

### DISCOVERY

Service discovery by name: cloud-level DNS, service-mesh level DNS, registration, client-side load balancing.

### OBSERVABILITY

Monitoring, metrics/health gathering, central tracing, APM, alarm management.  
Avoid "needle in a haystack" debugging.

**DEVOPS: SLAPS  
ROOF OF KUBERNETES**



**THIS BAD BOY CAN ORCHESTRATE  
SO MANY CONTAINERS IN IT**

# MICROSERVICES

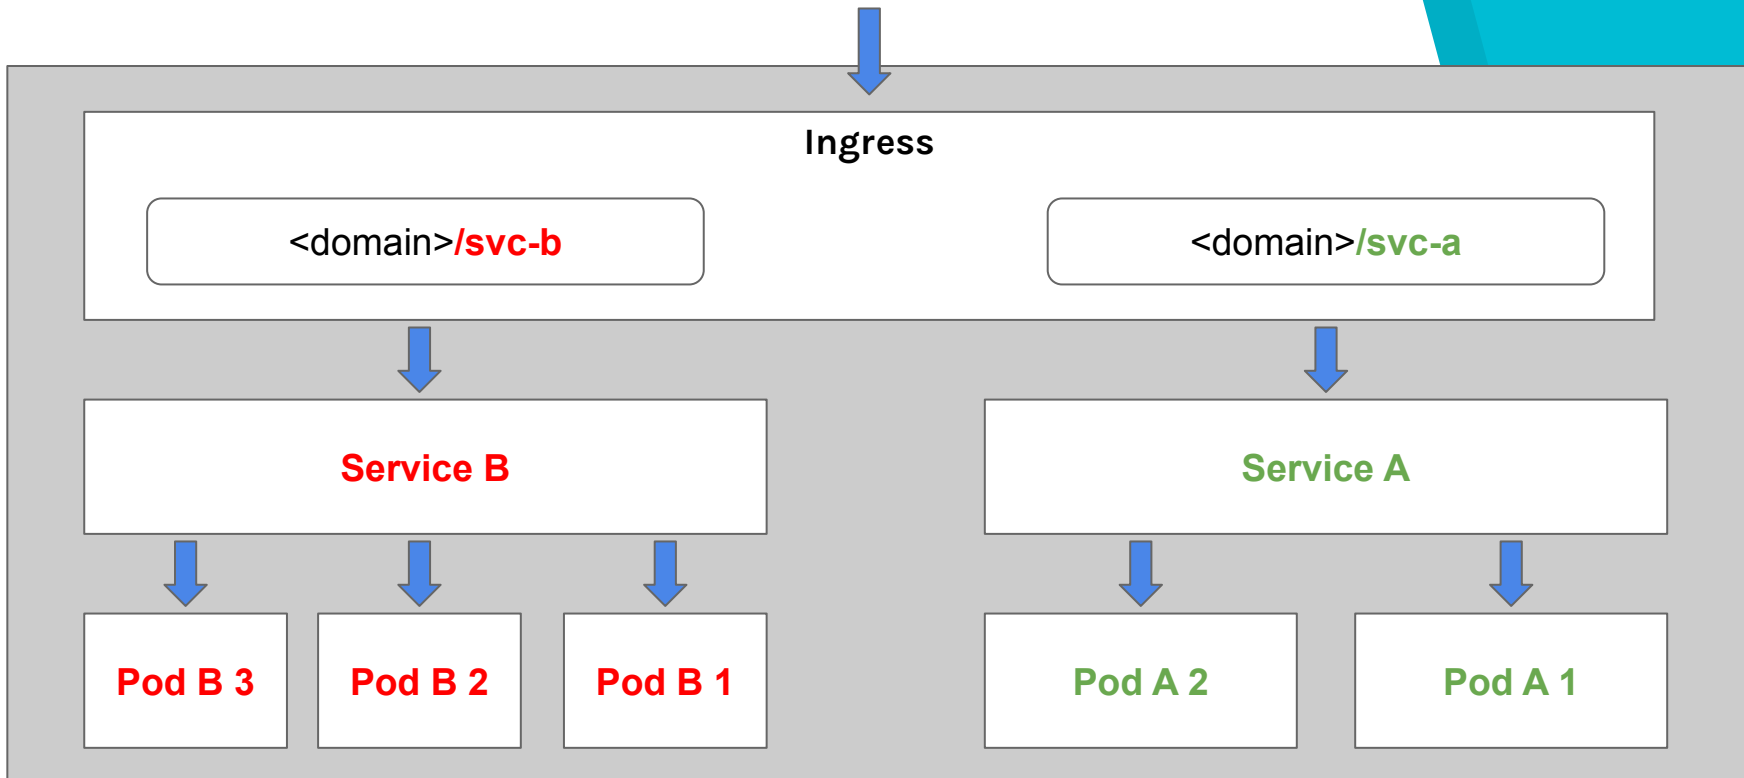
## Anatomy (Kubernetes)

- ▶ **Deployment**
  - ▷ Pod: Container (Image) + Resource limits
  - ▷ Replica set: count + deploy strategy
- ▶ **Service**
  - ▷ Deployment load balancer
  - ▷ Ingress traffic
- ▶ **HorizontalPodAutoScaler**
  - ▷ Deployment replica set horizontal scaling
- ▶ **ServiceMonitor**
  - ▷ Prometheus metrics exporter
- ▶ **Ingress (\*)**
  - ▷ Cluster gateway L7 router

# MICROSERVICES

## Anatomy (Kubernetes)

External Traffic



SHOW ME THE CODE!



**kubernetes**

# Deployment

```
# Deployment                                You, 18/01/2019 14:54 • Feature: k8s updated manifests
apiVersion: apps/v1beta2
kind: Deployment
metadata:
  name: '${ trigger.properties["application"] }'
  namespace: production
  labels:
    app: '${ trigger.properties["application"] }'
spec:
  replicas: 1
  revisionHistoryLimit: 2
  selector:
    matchLabels:
      app: '${ trigger.properties["application"] }'
  strategy:
    rollingUpdate:
      maxSurge: 1
      maxUnavailable: 0
    type: RollingUpdate
  template:
    metadata:
      labels:
        app: '${ trigger.properties["application"] }'
    spec:
      containers:
        - name: '${ trigger.properties["application"] }'
          image: '${ trigger.properties["imageUrl"] }'
          env:
            - name: NODE_ENV
              value: production
            - name: API_PORT
              value: '3000'
            - name: MONGODB_URI
              value: '${ trigger.properties["environment"]["production"]["MONGODB_URI"] }'
          livenessProbe:
            httpGet:
              path: /health
              port: 3000
            initialDelaySeconds: 30
            periodSeconds: 30
            successThreshold: 1
          ports:
            - containerPort: 3000
      resources:
        limits:
          cpu: 500m
```



# Service

```
# Service
apiVersion: v1
kind: Service
metadata:
  name: '${ trigger.properties["application"] }'
  namespace: production
  labels:
    app: '${ trigger.properties["application"] }'
spec:
  ports:
    - name: http
      port: 80
      protocol: TCP
      targetPort: 3000
  selector:
    app: '${ trigger.properties["application"] }'
  type: ClusterIP
```

# HPA

```
# HorizontalPodAutoscaler
apiVersion: autoscaling/v2beta1
kind: HorizontalPodAutoscaler
metadata:
  name: '${ trigger.properties["application"] }'
  namespace: production
spec:
  minReplicas: 2
  maxReplicas: 10
  scaleTargetRef:
    apiVersion: apps/v1beta2
    kind: Deployment
    name: '${ trigger.properties["application"] }'
  metrics:
    - resource:
        name: cpu
        targetAverageUtilization: 65
      type: Resource
    - resource:
        name: memory
        targetAverageUtilization: 65
      type: Resource
```

# Ingress

```
# Ingress
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: '${ trigger.properties["application"] }'
  namespace: production
  annotations:
    nginx.ingress.kubernetes.io/rewrite-target: /
  labels:
    app: '${ trigger.properties["application"] }'
spec:
  rules:
  - host: '${ trigger.properties["gatewayUrl"] }'
    http:
      paths:
      - path: '${ trigger.properties["servicePath"] }'
        backend:
          serviceName: '${ trigger.properties["application"] }'
          servicePort: 80
```

# SVC Monitor

```
# ServiceMonitor
apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
  name: '${ trigger.properties["application"] }'
  namespace: monitoring
  labels:
    app: '${ trigger.properties["application"] }'
    release: prometheus
spec:
  selector:
    matchLabels:
      app: '${ trigger.properties["application"] }'
  endpoints:
  - port: http
    path: /metrics
  namespaceSelector:
    matchNames:
      - production
```



A large teal graphic element consisting of a diagonal line that splits the page into a white upper-left section and a teal lower-right section. The teal section is a solid color, while the white section contains the main text.

# 2.

## CONTINUOUS DELIVERY

Concept  
Pipeline



“

Produce software in short cycles

Reliably released at any time

Building, testing and releasing faster  
and more frequently

Straightforward and repeatable  
deployment process

# PIPELINE

Code Push + Webhook



Lint



Test + Coverage



Quality Metrics Check



Build + Push Image



Deploy (Canary, Blue/Green, etc)


# 3.

## CANARY RELEASE




Concept  
Pipeline


Red/black  
(Blue/green)



Rolling red/black

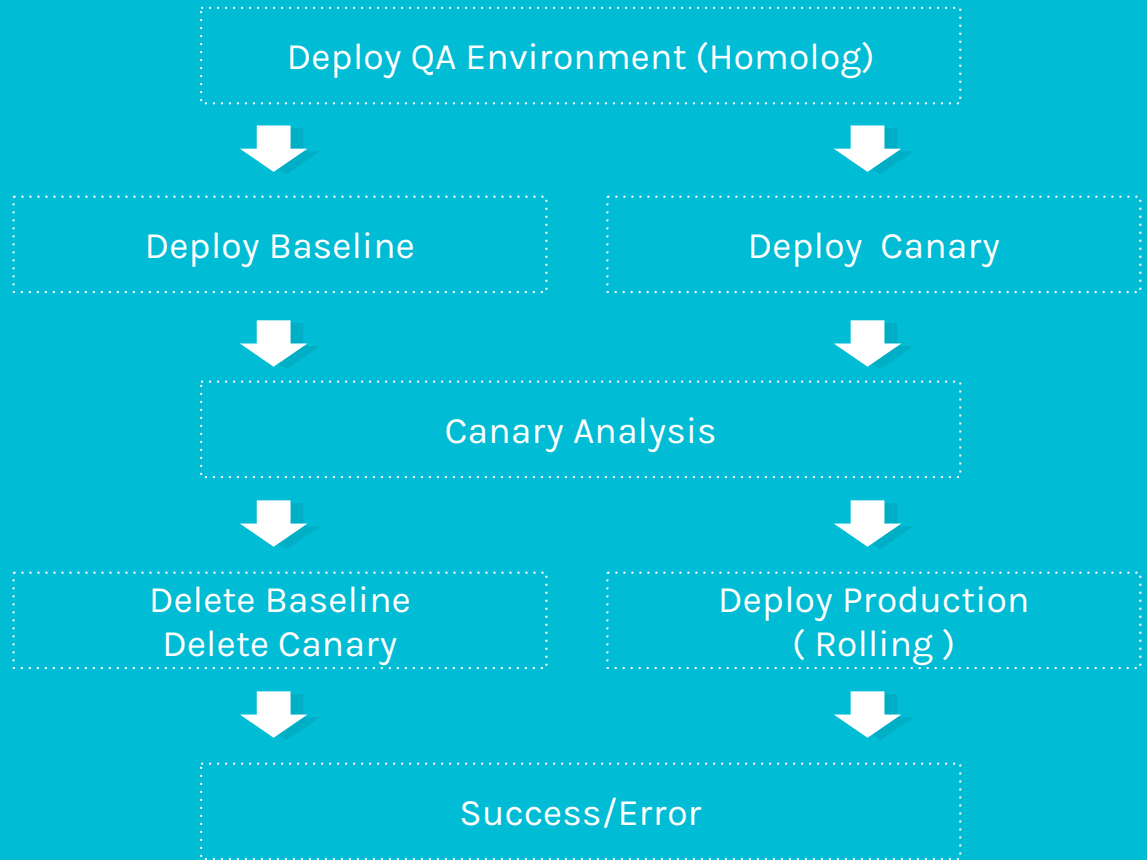


Canary







# PIPELINE





# 4.

## TOOLS

Topology  
Automation Server  
Integration Tests  
Code Quality  
Deploy  
Canary



Webhook



Trigger CD



Deploy



Canary Analysis



Code Quality Check  
(REST calls)



(LEEEROOOOOOOOOOY...)  
**JENKINS**

- ▶ Automation Server
- ▶ Pipeline
- ▶ Groovy DSL
- ▶ Plugins



(LEEEROOOOOOOOOOY...)  
**JENKINS**

- ▶ Kubernetes Plugin
- ▶ Custom Library
- ▶ Opinionated DSL



## Branches (4)

S	W	Name ↓	Last Success	Last Failure	Last Duration	Fav
		<a href="#">adjust-import-scripts</a>	1 mo 5 days - <a href="#">#2</a>	N/A	1 min 30 sec	 
		<a href="#">feature/livia-environment</a>	11 hr - <a href="#">#1</a>	N/A	1 min 8 sec	 
		<a href="#">master</a>	1 day 14 hr - <a href="#">#23</a>	N/A	2 min 38 sec	 
		<a href="#">postman-api-doc</a>	N/A	N/A	N/A	 




## Stage View

				Checkout	Install Dependencies	Lint	Tests	Check quality: Scan	Check quality: Quality Gate	Build and publish image	Deploy properties
Average stage times: (Average full run time: ~2min 32s)				2s	22s	4s	7s	21s	6s	57s	6s
<a href="#">#23</a>	Jan 16 11:58	1 commit		2s	21s	5s	7s	22s	6s	1min 0s	7s
<a href="#">#22</a>	Jan 15 20:18	No Changes		2s	26s	4s	6s	19s	4s	51s	6s
<a href="#">#21</a>	Jan 15 18:33	No Changes		2s	21s	5s	7s	20s	5s	50s	7s




App Test Pod

App Container



DB Container



Cache Container



**SHOW ME THE CODE!**





# Pipeline

```
@Library('livia-tools-jenkins-pipeline-lib@master') _

def serviceName = "tools-featureflag-api"
def servicePath = "/featureflag"
def registryRepository = "${serviceName}-prd"
def containers = [
    containerTemplate(
        name: 'mongo',
        image: 'mongo:latest',
        ttyEnabled: true,
        ports: [portMapping(name: 'mongodb', containerPort: 27017, hostPort: 27017)]
    )
]
def testEnv = [MONGODB_URI: "mongodb://localhost:27017/${serviceName}"]
def homologEnv = [
    NODE_ENV: 'production',
    MONGODB_URI: 'mongodb://****/featureflag'
]
def productionEnv = [
    NODE_ENV: 'production',
    MONGODB_URI: 'mongodb://****/featureflag?replicaSet=****'
]

nodeJSApi(
    name: serviceName,
    servicePath: servicePath,
    branch: env.BRANCH_NAME,
    nodeImage: 'node:dubnium',
    testEnv: testEnv,
    homologEnv: homologEnv,
    productionEnv: productionEnv,
    containers: containers,
    registryRepository: registryRepository,
    canaryEnabled: true
)
```

# DSL

```
def call(Map args) {
  // -- arguments --
  def name = args.name
  def branch = args.branch
  def servicePath = args.servicePath ?: args.name

  def nodeImage = args.nodeImage ?: 'node:dubnium'

  def testEnv = args.testEnv ?: [:]
  def homologEnv = args.homologEnv ?: [:]
  def productionEnv = args.productionEnv ?: [:]

  def containers = args.containers ?: []
  def volumes = args.volumes ?: []

  def registryRepository = args.registryRepository ?: name
  def canaryEnabled = args.canaryEnabled ?: false
  def masterBranch = branch == 'master'

  // -- global constants --
  def DEFAULT_CONTAINERS = [
    containerTemplate(name: 'node', image: nodeImage, ttyEnabled: true),
    containerTemplate(name: 'docker', image: 'docker', ttyEnabled: true),
    containerTemplate(name: 'kubect1', image: 'lachlanevenson/k8s-kubect1:v1.10.5', command: 'cat', ttyEnabled: true)
  ]
  def DEFAULT_VOLUMES = [hostPathVolume(hostPath: '/var/run/docker.sock', mountPath: '/var/run/docker.sock')]
  def REGISTRY_CONFIG = [
    host: '****.amazonaws.com',
    url: "https://****.amazonaws.com",
    credentials: "ecr:****"
  ]
}
```

```
// -- execution --
def label = "job-UUID.randomUUID().toString()"
def podContainers = DEFAULT_CONTAINERS + containers
def podVolumes = DEFAULT_VOLUMES + volumes

podTemplate(label: label, containers: podContainers, volumes: podVolumes, serviceAccount: 'jenkins') {
  node(label) {
    notifyStatus {
      // ** CLONE REPO **
      stage('Checkout') {
        checkout scm
      }

      container('node') {
        // ** INSTALL DEPENDENCIES **
        stage('Install Dependencies') {
          sh 'npm install'
        }
        // ** LINT CHECK **
        stage('Lint') {
          sh 'npm run lint'
        }
        // ** UNIT AND INTEGRATION TESTS **
        stage('Tests') {
          def testEnvList = testEnv.collect({ key, value -> "${key}=${value}" })
          withEnv(testEnvList) {
            sh 'npm run test:ci:cover'
          }
        }
      }
    }
  }
}
```

```
if (masterBranch) {
  // ** SONAR SCANNER **
  stage('Check quality: Scan') {
    sh 'npm run sonar:scanner'
  }
}

if (masterBranch) {
  // ** CODE QUALITY CHECKS **
  stage('Check quality: Quality Gate') {
    def versionFolders = listSubFolders "api/v*"
    def qualityGateStatus = getQualityGateStatus versionFolders
    if (!qualityGateStatus) {
      error "Failed because at least one version did not pass the Quality Gate"
    } else {
      echo "Passed Quality Gate check"
    }
  }
}

// ** BUILD AND PUBLISH **
def tag = getGitCommit()
def imageName = "${registryRepository}:${tag}"
container('docker') {
  stage('Build and publish image') {
    docker.withRegistry(REGISTRY_CONFIG.url, REGISTRY_CONFIG.credentials) {
      docker
      .build(imageName)
      .push()
    }
  }
}
```

```
stage("Deploy properties") {
  def imageUrl = "${REGISTRY_CONFIG.host}/${imageName}"

  def deployData = [
    application: name,
    servicePath: servicePath,
    imageUrl: imageUrl,
    environment: [
      homolog: homologEnv,
      production: productionEnv,
    ]
  ]

  if (canaryEnabled) {
    echo 'Canary enabled: Building canary config'
    container('kubect1') {
      def currentProductionImageUrl = getRunningPodImage app: name, namespace: 'production'
      def currentProductionState = getRunningPodState app: name, namespace: 'production'
      deployData.put 'canary', [
        baselineState: currentProductionState,
        baselineImageUrl: currentProductionImageUrl,
        canaryImageUrl: imageUrl
      ]
    }
  }

  archiveDeployTriggerYaml name: 'deploy.yaml', data: deployData
}
```

# Running Pipeline Pod

```
Rodrigos-MacBook-Pro-2:~ rodrigobotti$ kubectl get pods -n jenkins -w
```

NAME	READY	STATUS	RESTARTS	AGE
jenkins-7669df5fd8-wvzpl	1/1	Running	0	29d
job-cedb6d14-1c3f-4caf-8d82-b009d5443f82-krnjg-fcmm3	0/5	Pending	0	1s
job-cedb6d14-1c3f-4caf-8d82-b009d5443f82-krnjg-fcmm3	0/5	Pending	0	1s
job-cedb6d14-1c3f-4caf-8d82-b009d5443f82-krnjg-fcmm3	0/5	ContainerCreating	0	1s
job-cedb6d14-1c3f-4caf-8d82-b009d5443f82-krnjg-fcmm3	5/5	Running	0	3s

# QUALITY CHECKS

## SONARQUBE

- ▶ Quality Metrics
  - ▷ Code smells
  - ▷ Bugs
  - ▷ Vulnerabilities
  - ▷ Code coverage
- ▶ Quality Profiles
- ▶ Quality Gate



**Conditions**


Only project measures are checked against thresholds. Sub-projects, directories and files are ignored. [More](#)

Metric	Over Leak Period	Operator	Warning	Error		
Blocker Issues	<input type="checkbox"/>	is greater than	<input type="text"/>	0	Update	Delete
Bugs	<input type="checkbox"/>	is greater than	<input type="text"/>	0	Update	Delete
Code Smells	<input type="checkbox"/>	is greater than	<input type="text"/>	0	Update	Delete
Confirmed Issues	<input type="checkbox"/>	is less than	<input type="text"/>	0	Update	Delete
Coverage	<input type="checkbox"/>	is less than	95	92	Update	Delete
Critical Issues	<input type="checkbox"/>	is less than	<input type="text"/>	0	Update	Delete
Duplicated Blocks	<input type="checkbox"/>	is greater than	<input type="text"/>	0	Update	Delete
Maintainability Rating	Never	is worse than	<input type="text"/>	A X	Update	Delete
Major Issues	<input type="checkbox"/>	is greater than	<input type="text"/>	0	Update	Delete
Minor Issues	<input type="checkbox"/>	is less than	<input type="text"/>	0	Update	Delete
Open Issues	<input type="checkbox"/>	is less than	<input type="text"/>	0	Update	Delete
Reliability Rating	Never	is worse than	<input type="text"/>	A X	Update	Delete
Security Rating	Never	is worse than	<input type="text"/>	A X	Update	Delete
Vulnerabilities	<input type="checkbox"/>	is greater than	<input type="text"/>	0	Update	Delete




Quality Gate Passed


Bugs 🔗 Vulnerabilities 🔗



Code Smells 🔗



Coverage 🔗



Duplications 🔗



XS 980  
Lines of Code  
JavaScript ■ 980

No tags ▾

Activity

January 16, 2019

1.0.0

January 15, 2019

Project Analyzed

December 21, 2018

Project Analyzed

[Show More](#)

Quality Gate

(Default) [Microservices](#)

Quality Profiles

(JavaScript) [Sonar way Recommended](#)

Key


tools:featureflag:api:v1

# MONITORING

## PROMETHEUS (OPERATOR)

- ▶ Monitoring
- ▶ Time-series database
- ▶ PromQL
- ▶ Alerting





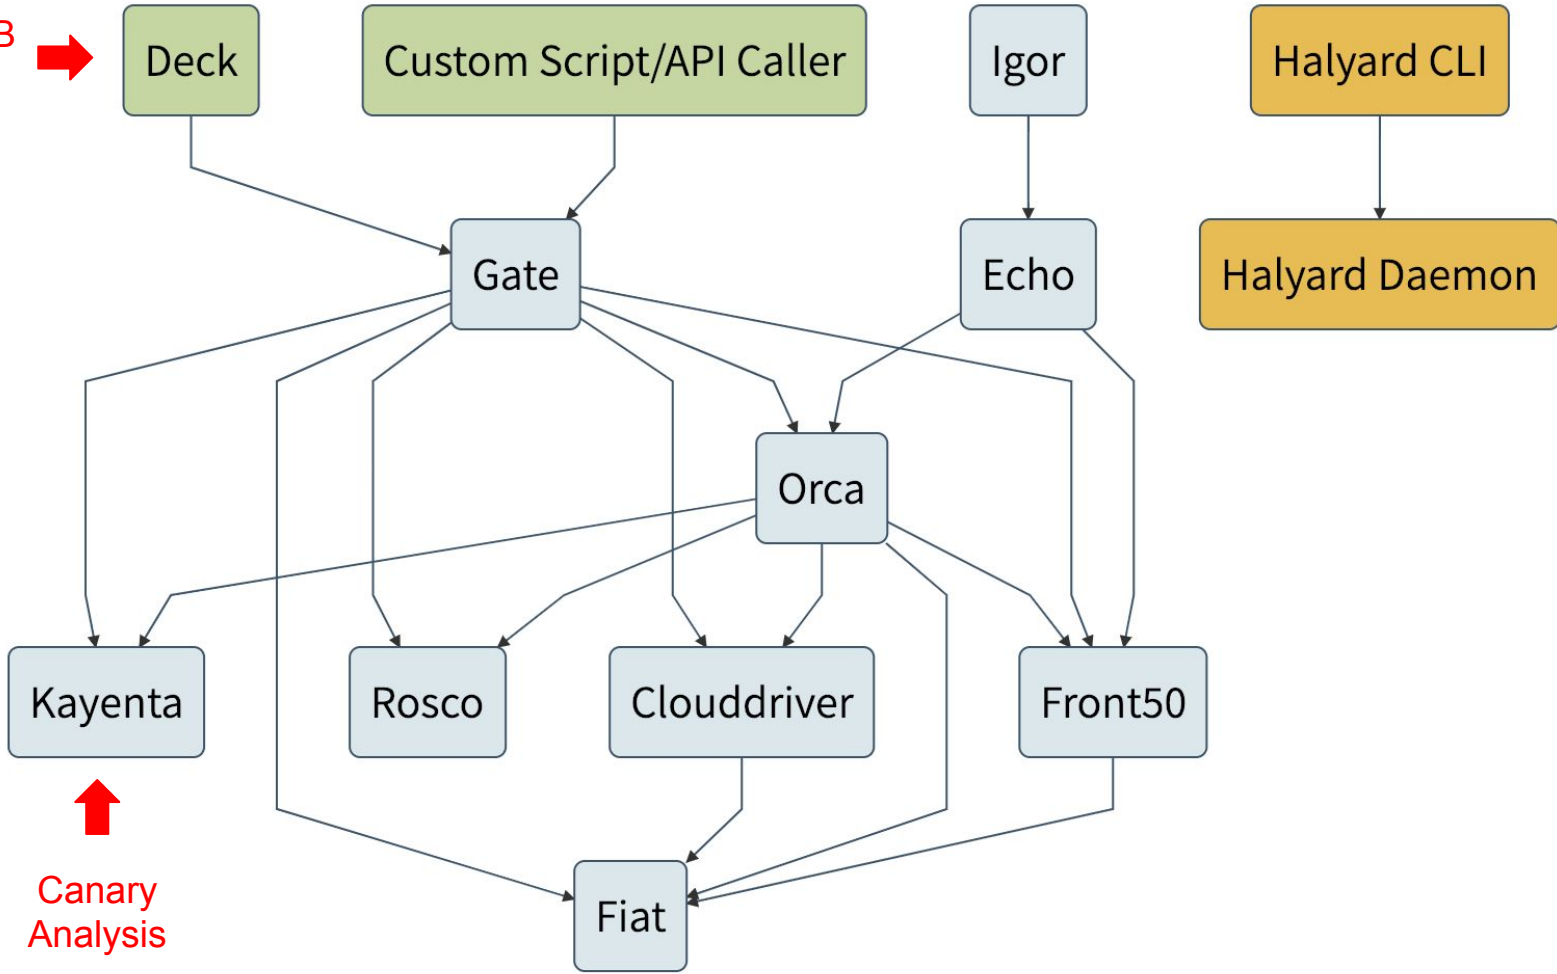
# DELIVERY

## SPINNAKER

- ▶ CD Pipeline
- ▶ Canary Release
- ▶ Multi-cloud
  - ▷ Kubernetes
- ▶ Abstracts infrastructure elements (\*)
- ▶ Netflix



WEB UI →



↑  
Canary Analysis

**YO DAWG, HEARD YOU LIKE SERVICES**

**N**

**SO WE PUT SERVICES IN A SERVICE SO  
YOU CAN DEPLOY SERVICES USING A SERVICE**






# CONCEPTS/ABSTRACTIONS

## SPINNAKER

- ▶ Application
- ▶ Cluster == **Namespace**
- ▶ Server Group == **Deployment**
- ▶ Load Balancer == **Service**





## HOMOLOG

 **V018:** 454706396284.dkr.ecr.sa-east-1.amazonaws.com/tools-featureflag-api-prd:644996926ea265c9a4f2ec140b7335c732f7de05   1 ▲ : 100%



Instance▼	Launch Time▼	Zone▼	Provider▼
▲ pod tools-featureflag-api-fcc - 8957c4-h48nz		homolog	Up

 **V017:** 454706396284.dkr.ecr.sa-east-1.amazonaws.com/tools-featureflag-api-prd:23b31addccb2aefad2a1eba7a6dc2bd164400e16  


## PRODUCTION

 **V022:** 454706396284.dkr.ecr.sa-east-1.amazonaws.com/tools-featureflag-api-prd:23b31addccb2aefad2a1eba7a6dc2bd164400e16   2 ▲ : 100%

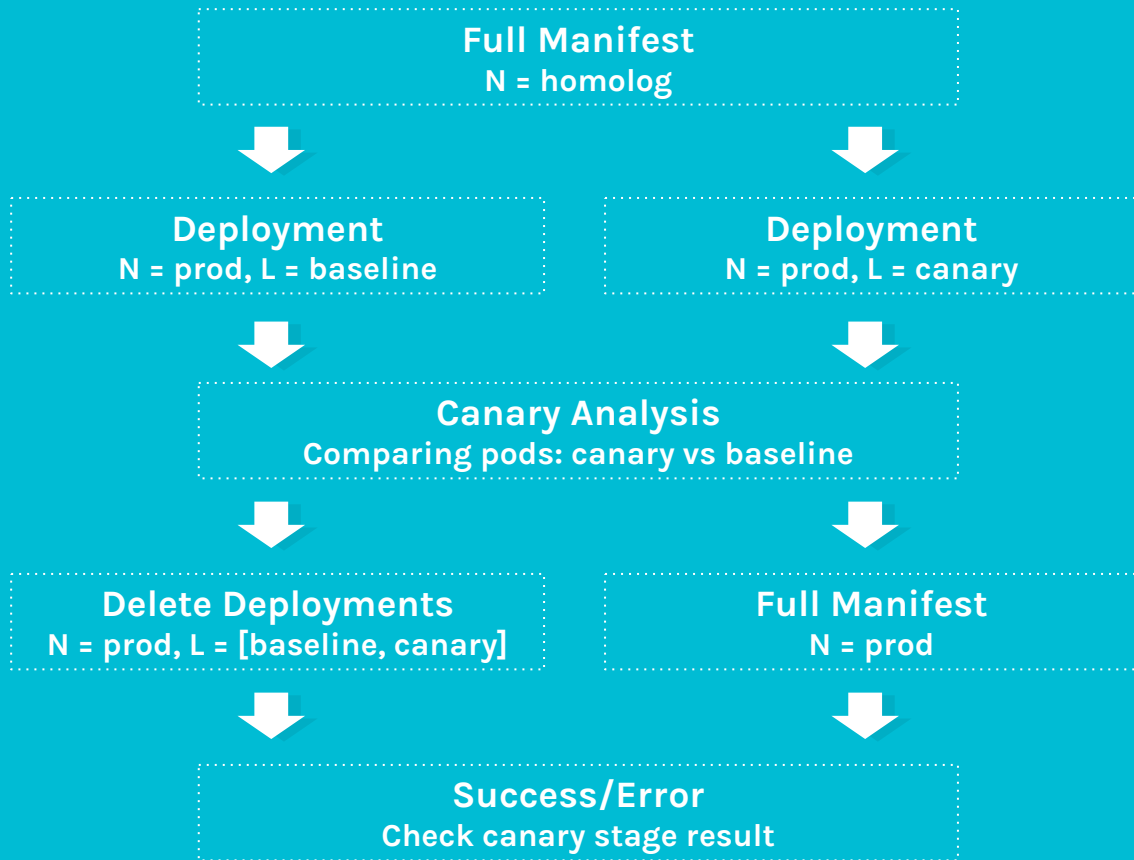
Instance▼	Launch Time▼	Zone▼	Provider▼
▲ pod tools-featureflag-api-65 - 4f8c5669-f4zhn		production	Up
▲ pod tools-featureflag-api-65 - 4f8c5669-pzdrq		production	Up

 **V021:** 454706396284.dkr.ecr.sa-east-1.amazonaws.com/tools-featureflag-api-prd:d4f217e2623d1920564cc4c980e12eddb3e0635e  

# PIPELINE



# PIPELINE





# Pipeline Overview


## deploy tools-featureflag-api

Permalink 

 Create

 Configure 

Pipeline Actions 



 Add stage

 Copy an existing stage

# Stage: Configuration

## Automated Triggers

Type	Jenkins  Listens to a Jenkins job	
Master	jenkins-aws 	
Job	tools-featureflag-api/job/master 	
Property File 	deploy.yml	
<input checked="" type="checkbox"/> Trigger Enabled		

```
application: tools-featureflag-api
servicePath: /featureflag
imageUrl: '****.amazonaws.com/tools-featureflag-api-prd:644996926ea265c9a4f2ec140b7335c732f7de05'
environment:
  homolog:
    NODE_ENV: production
    MONGODB_URI: mongodb://****/featureflag
  production:
    NODE_ENV: production
    MONGODB_URI: mongodb://****/featureflag?replicaSet=****
canary:
  baselineState:
    imageUrl: '****.amazonaws.com/tools-featureflag-api-prd:23b31adddcb2aefad2a1eba7a6dc2bd164400e16'
    environment:
      NODE_ENV: production
      API_PORT: '3000'
      MONGODB_URI: mongodb://****/featureflag?replicaSet=****
  baselineImageUrl: '****.amazonaws.com/tools-featureflag-api-prd:23b31adddcb2aefad2a1eba7a6dc2bd164400e16'
  canaryImageUrl: '****.amazonaws.com/tools-featureflag-api-prd:644996926ea265c9a4f2ec140b7335c732f7de05'
```

# Stage: Canary Analysis

## Canary Config

Analysis Type <sup>?</sup> Real Time

### Analysis Config

Config Name

Delay <sup>?</sup>  minutes before starting analysis

Interval <sup>?</sup>  minutes

Lookback Type <sup>?</sup> Growing

### Metric Scope

Baseline <sup>?</sup>

Baseline Location <sup>?</sup>

Canary <sup>?</sup>

Canary Location <sup>?</sup>

Step  seconds

Lifetime <sup>?</sup>  hours  minutes

Resource Type

Extended Params <sup>?</sup>

Key	Value
-----	-------

### Scoring Thresholds

Marginal <sup>?</sup>

Pass <sup>?</sup>

### Advanced Settings

Metrics Account <sup>?</sup>

Storage Account <sup>?</sup>

Scope Name

Configuration Name

Description

## METRICS

ALL

**RESPONSE\_TIME**

Add Group

METRIC NAME

request\_duration

Add Metric

GROUPS

response\_time



## FILTER TEMPLATES

TEMPLATE NAME

filter\_application

Add Template



## SCORING

Thresholds

Marginal ⓘ

Pass ⓘ

Judge

NetflixACAJudge-v1.0

Metric Group Weights ⓘ

response\_time



# Canary Config Metric

## Configure Metric

<b>Group</b>	<input type="text" value="response_time"/>
<b>Name</b>	<input type="text" value="request_duration"/>
<b>Fail on</b>	<input checked="" type="radio"/> Increase <input type="radio"/> Decrease <input type="radio"/> Either
<b>Criticality</b>	<input checked="" type="checkbox"/> Fail the canary if this metric fails
<b>NaN Strategy</b>	<input checked="" type="radio"/> Default (remove) <input type="radio"/> Replace with zero <input type="radio"/> Remove
<b>Filter Template</b>	<input type="text" value="filter_application"/>
<b>Scope Name</b>	<input type="text" value="default"/>
<b>Metric Name</b>	<input type="text" value="http_request_duration_ms_sum"/>
<b>Label Bindings</b>	<input type="text" value="Add new"/>
<b>Group By</b>	<input type="text" value="Add new"/>

Cancel

OK

# Canary Config Filter

### Edit Template

<b>Name</b>	<input type="text" value="filter_application"/>
<b>Template</b>	<input \${="" scope="" type="text" value="pod=~" }.*",namespace="\${ location }"/>

---

# Canary Analysis Report: Success

Delivery

PIPELINES 1 CANARY CONFIGS CANARY REPORTS

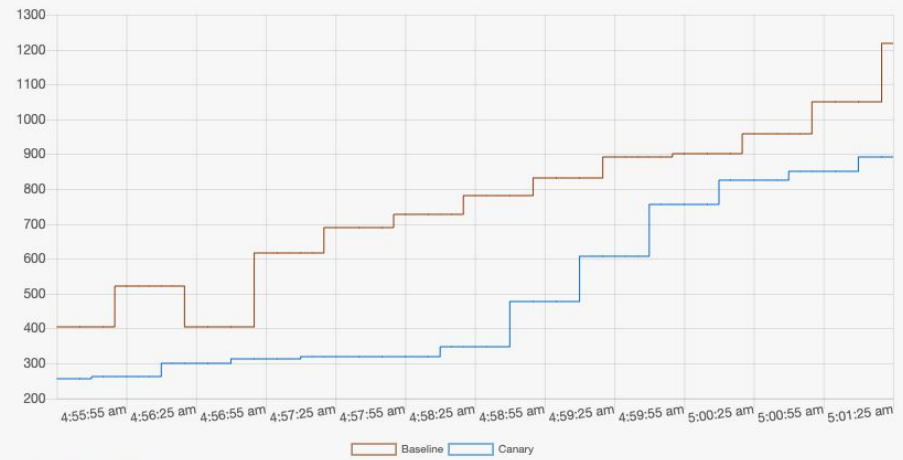
response **100**  
PASS

BASILINE SCOPE	LOCATION	CANARY SCOPE	LOCATION	TIME START	END	STEP	THRESHOLD MARGINAL	PASS	SOURCE
tools-featureflag-api-baseline	production	tools-featureflag-api-canary	production	2019-02-04 04:55:25 PST	2019-02-04 05:01:25 PST	0.08333333333333333 mins	70	90	Report Metrics

ALL

RESPONSE\_TIME

METRIC NAME	RESULT
request_duration	Pass



GRAPH: [Amplitude vs. Time](#) Histogram

NAME: request\_duration  
QUERY: http\_request\_duration\_ms\_sum

	COUNT	AVG	MAX	MIN
Baseline	73	749.23	1218	404
Canary	73	501.66	893	254.5

# Canary Analysis Report: Failed

Deployment: tools-featureflag-api

Delivery PIPELINES CANARY CONFIGS CANARY REPORTS

response **0 FAIL**

BASILINE SCOPE	LOCATION	CANARY SCOPE	LOCATION	TIME START	END	STEP	THRESHOLD MARGINAL	PASS	SOURCE
tools-featureflag-api-baseline	production	tools-featureflag-api-canary	production	2019-04-23 22:45:37 PDT	2019-04-23 22:47:37 PDT	0.08333333333333333 mins	70	90	Report Metrics

**ALL** **RESPONSE\_TIME**

METRIC NAME: request\_duration RESULT: Nodata

GRAPH: [Amplitude vs. Time](#) Histogram

NAME: request\_duration

QUERY: http\_request\_duration\_ms\_sum

	COUNT	AVG	MAX	MIN
Baseline	25	66.72	100	35
Canary	0	0	0	0

# Stage: Check

## Edit Precondition ✕

**Check** Expression ⌵

**Expression** ⓘ `${ #stage('Canary Analysis')['status'].toString() == 'SUCCEEDED' }`

**Fail Pipeline** ⓘ

Cancel Update



**5.**

**FUTURE**

# P.O.C.

## kubernetes-cluster-sa-east

### namespace: spin

Spinnaker:

Deck  
Gate  
Kayenta

...

\* in single box for  
diagram simplicity

### namespace: jenkins

Jenkins master

Jenkins agent

Jenkins job slave

Jenkins job slave

### namespace: production

feature toggle

profiles

labs-properties

## k8s-cd-cluster-us-east

namespace: spin

Spinnaker

namespace: jenkins

Jenkins

## k8s-app-cluster-us-east

namespace: production

migrated

## GKE

namespace: default

migrating




# ONGOING

- ▶ **Migrating CIs to Jenkins**
  - ▷ Wercker
  - ▷ **\* CircleCI \***
  - ▷ Saves \$ + Higher concurrency
  - ▷ **Standardizes** pipeline and **quality** measures
- ▶ **Jenkins pushing to both registries**
  - ▷ ECR
  - ▷ GCR
- ▶ **Spinnaker aware of both clusters**
- ▶ **Kayenta metrics**
  - ▷ Response time
  - ▷ Error rates
  - ▷ Resource consumption

# FUTURE

- ▶ Canary Analysis config improvement
  - ▷ What metrics?
  - ▷ Groups weights
  - ▷ Length/Windows
- ▶ Service Mesh?
- ▶ Manifests from SCM
- ▶ Pipeline in SCM
- ▶ Canary Analysis with business metrics
  - ▷ Custom collector?
- ▶ Automatic feature toggle management



**LETS DEPLOY**



**DIRECT TO PRODUCTION**

6.

Q&A





**7.**

**"THANK YOU"  
NOTES**



**rivendel**



Reúna todos os  
seus exames em  
um só lugar

<https://livia.saude.com.br>






# THE DEVELOPER'S CONFERENCE







## Contact

rodrigo.botti@gmail.com

rodrigo.botti@nexadigital.com.br

<https://www.linkedin.com/in/rodrigo-botti/>

