@MELANIECEBULA / SEPT 2018 / ONS EUROPE

Services for All

How to Empower Engineers (with Kubernetes)



- 1. A brief history of configuration at Airbnb
- 2. Generating Kubernetes files
- 3. Creating a kubectl wrapper
- 4. Extending and customizing Kubernetes
- 5. Beyond Kubernetes



Start planning your trip



Homes



Experiences



Restaurants

Experiences near your home in Amsterdam



BOAT RIDE Amsterdam Experience Cruise \$59 per person 4.89 ***** 836



FOOD WALK The all Dutch food & history tour \$91 per person 4.91 ***** 196



PHOTO WALK Capture the city & you on a photo walk \$35 per person 4.83 ***** 132

Show all (118) >

Where to stay



PRIVATE ROOM · 1 BED LUXURY INDEPENDENT STUDIO on SHIP : free bikes! \$162 per night · Free cancellation ★★★★★ 334 · Superhost



PRIVATE ROOM · 2 BEDS Rebel - Private Room \$126 per night · Free cancellation ★★★★★ 551 · Superhost



PRIVATE ROOM · 1 BED Bed & Boat, apartment on houseboat. Free bikes. \$145 per night · Free cancellation ★★★★ 328 · Superhost

Show all (2000+) >

BOAT RIDE **Romantic Waters Boattour** \$35 per person 4.87 **** 231



BIKE RIDE Jewish tour & visit to AnneFrank house \$43 per person 4.79 ***** 24



BOAT RIDE Early morning- Canals all to ourselves \$42 per person 4.89 **** 565





PRIVATE ROOM · 1 BED Private Attic Studio/Roofterrace \$90 per night · Free cancellation ★★★★ 370 · Superhost



PRIVATE ROOM · 1 BED Experience a houseboat in Amsterdam \$145 per night · Free cancellation ★★★★★ 247 · Superhost



PRIVATE ROOM · 1 BED Authentic houseboat with privacy and comfort \$128 per night · Free cancellation ★★★★★ 373 · Superhost







What is airbnb?

AN ONLINE MARKETPLACE FOR SHARING HOMES AND EXPERIENCES

81k cities

@MELANIECEBULA

191 +countries

5 mil homes



Whoam I?



A BRIEF HISTORY





	1000	
	750	
ENGINEERING TEAM	500	
	250	
	0	009 2010



Why Microservices? **SCALING CONTINUOUS DELIVERY**









@MELANIECEBULA

120,000 production deploys per year

Splitting the Monolith

Monolith



Splitting the Monolith



@MELANIECEBULA

Standardize Service Creation

NO SNOWFLAKES



Standardize Service Creation

OKAY BUT HOW?

CONFIGURATION IS KEY



Evolution of Configuration Management



Manually configuring boxes

Automating configuration of @MELANIECEBULA







applications with Chef

Automating configuration and orchestration of containerized applications with Kubernetes

KUBERNETES OUT-OF-THE-BOX



What Makes Kubernetes Awesome

- immutable, reproducible containers
- designed for a microservices architecture
- human-readable format
- declarative
- efficient scheduling
- extensible API

What Makes **Kubernetes NOT** Awesome

- complex configuration and concepts
- tooling is not developer friendly
- open issues

significant set up cost

GENERATING KUBERNETES FILES



kubernetes config files



kubernetes



kubernetes **IS REPETITIVE**









bonk-canary/ bonk-development/ bonk-production/ bonk-staging/

/Users/melanie_cebula/onetouch-codelabs/projects/bonk/generated-apps/bonk/

```
bonk-production-admin-role-binding.yml
bonk-production-databag-bonk-configmap.yml
bonk-production-deployment.yml
bonk-production-mini-announcer-configmap.yml
bonk-production-service-config-map-configmap.yml
bonk-production-service.yml
bonk-production-synapse-configmap.yml
bonk-production-zoned-key-configmap.yml
```



@MELANIECEBULA

What are the basic attributes of my project?



- # under _infra/kube/kube-gen.yml
- version: 5.1.0
- project:
 - name: bonk
- environments:
 - production:
 - params:
 - replicas: 10
 - port: 6585
 - staging:
 - params:
 - replicas: 2
 - port: 6586
- users:

melanie_cebula





- # under _infra/kube/kube-gen.yml
- version: 5.1.0
- project:
 - name: bonk
- environments:
 - production:
 - params:
 - replicas: 10
 - port: 6585
 - staging:
 - params:
 - replicas: 2
 - port: 6586
- users:

melanie_cebula

becomes an admin role binding file!



kube-gen GENERATED ADMIN ROLEBINDING FILE



```
# under generated/bonk/bonk-production/admin-role-
binding.yml
```

apiVersion: rbac.authorization.k8s.io/v1

```
kind: RoleBinding
```

metadata:

```
creationTimestamp: null
```

```
name: bonk-production-admin-role-binding
```

```
namespace: bonk-production
```

roleRef:

apiGroup: rbac.authorization.k8s.io

kind: ClusterRole

name: admin

subjects:

- kind: User

name: melanie_cebula





under _infra/kube/kube-gen.yml version: 5.1.0 project: name: bonk environments can be accessed by all the environments: other files! params: replicas: 10 port: 6585 params: replicas: 2 port: 6586 users:

melanie_cebula



kube-gen **TEMPLATE VARIABLES**



different concepts map to different yaml files under _infra/kube

environment params defined in the project are accessed with go templating:

ex: {{ .Env.Params.replicas }}

```
# under _infra/kube/apps/bonk.yml
 deployment:
   replicas: {{ .Env.Params.replicas }}
   strategy:
      rollingUpdate
```



kube-gen **TEMPLATE VARIABLES**



{{ }} signals go templating

kube-gen will generate all the files for the different environments, and replace each .Env.Params.replicas with the appropriate value from the project file

ex: staging has 2 replicas, production has 10 replicas

```
# under _infra/kube/apps/bonk.yml
  deployment:
                                              from the
   replicas: {{ .Env.Params.replicas }}
                                             project file!
   strategy:
      rollingUpdate
```





kube-gen **APP FILES**



@MELANIECEBULA

What kind of workload?

example workloads



kube-gen APP FILES



under _infra/kube/apps/bonk-web.yml
workload:

deployment:

{{ if eq .Env.Name "development" }}
strategy:
rollingUpdate:
maxUnavailable: 1

{{ else }}

autoscaling:

minReplicas: {{ .Env.Params.minReplicas }}

maxReplicas: {{ .Env.Params.maxReplicas }}

{{ end }}



kube-gen COMPONENTS



@MELANIECEBULA

Which shared components to use?

example components



kube-gen **COMPONENTS**

nginx component



common patterns are abstracted into a component apps enable components component yaml merged into project may require params to be set may set default params

kube-gen **CONTAINER FILES**



@MELANIECEBULA

What does the container need?

kube-gen **CONTAINER FILES**



- image to use
- command and args to run
- ENV variables to pass in
- resource requests and limits for container
- file and volume mounts used by this container

kube-gen **OTHER FILES**



These can be mounted into your containers

kube-gen **OTHER FILES**



The main Dockerfile for the project (installs needed dependencies)

KUBECTL WRAPPER



kubernetes config files



kubernetes

kubernetes config files



kubect IS VERBOSE

k too KUBECTL WRAPPER





ktool The All purpose cli wrapper



ktoo **USES ENV VARS**

• Runs in the project home directory: \$ cd /path/to/bonk

\$ k status

• Environment variables for arguments:

- \$ k status ENV=staging
- \$ export ENV=staging
- \$ k status
- Prints the command that it will execute:
 - \$ k status ENV=staging

kubectl get pods --namespace=bonk-staging

ktoo SIMPLIFIES BUILD AND DEPLOYS

- files
- tags

• k generate transforms kube-gen files to kubernetes

• k build performs docker build and docker push with

• k deploy creates namespace, applies/replaces kubernetes files, sleeps and checks deployment status • can chain commands; ex: k all

ktoo A DEBUGGING TOOL

- specify particular pod, specific container:
- proxy
- k diagnose: shows status and logs for each failing container of your pod, and shows failure events

- defaults to random pod, main container:
 - \$ k ssh ENV=staging
 - \$ k logs ENV=staging POD=... CONTAINER=statsd-

EXTENDING KUBERNETES



WITH CUSTOM CONTROLLERS

admission controller:

- require project ownership (via annotations)
- prevent namespace creation under incorrect cluster —
- _
- - automatically rotate old pods
 - rotates over-provisioned availability zones —

prevent deployment of objects with known security holes

deployment pruner:

WITH GRACEFUL TERMINATION

- it is a common anti-pattern for clients of a service to not retry on connection failures
- as a k8s service deploys, requests may be sent to terminating pods and cause a spike in 500s
- use graceful termination to mitigate this

WITH GRACEFUL TERMINATION

under _infra/kube/apps/bonk.yml # wait up to 180 seconds before "kill -9" the pods terminationGracePeriodSeconds: 180

- # under _infra/kube/containers/container.yml
- # wait up to 120 seconds before shutting down
- lifecycle:
 - preStop:
 - exec:
 - command:
 - /bin/sleep
 - "120"

WITH GRACEFUL TERMINATION

under _infra/kube/apps/bonk.yml # wait up to 180 seconds before "kill -9" the pods terminationGracePeriodSeconds: 180

preStop:

under _infra/kube/containers/container.yml # wait up to 120 seconds before shutting down lifecycle:

exec:

command:

- /bin/sleep

- "120"

this gives our service discovery container time to mark this pod as unhealthy

WITH GRACEFUL TERMINATION

under _infra/kube/apps/bonk.yml # wait up to 180 seconds before "kill -9" the pods terminationGracePeriodSeconds: 180

- # under _infra/kube/containers/container.yml # wait up to 120 seconds before shutting down lifecycle: preStop:
 - exec:
 - command:
 - /bin/sleep
 - "120"

you can do more sophisticated server shutdown here too!

WITH GRACEFUL TERMINATION



- # under _infra/kube/containers/container.yml
 - # wait up to 120 seconds before shutting down
- lifecycle:
 - preStop:
 - exec:
 - command:
 - /bin/sleep
 - "120"





BEYOND KUBERNETES



Creating a new service (before)



Everything about a service should be in one place, and managed with one process.

Configuration LIVES IN ONE PLACE

/Users/melanie_cebula/bonk/

- _infra/
 - ▶ ci/
 - docs/
 - keys/
 - kube/
 - secrets/
 airlab.yml
 - aws.yml
 - deployboard.yml dyno.yml project.yml
- app/
- bin/
- config/
- ▶ db/
- ▶ lib/
- log/
- > public/
- spec/
- ▶ tmp/
- vendor/
 config.ru
- Comfile
- Gemfile
- Gemfile.lock
- Rakefile
- README.md
- unicorn.rb

Everything you need to know about your service can be found in one place

- · All configuration lives in _infra alongside project code
- Edit code and configuration with one pull request
- Easy to add new configuration
- Configuration statically validated as part of Cl



Configuration LIVES IN ONE PLACE

/Users/melanie_cebula/bonk/

- 🗕 _infra/
 - ▶ ci/
 - docs/
 - keys/
 - kube/
 - secrets/ airlab.yml
 - aws.yml
 - deployboard.yml dyno.yml project.yml
- ▶ app/
- bin/
- config/
- db/
- lib/
- ▶ log/
- > public/
- spec/
- tmp/
- vendor/ config.ru
- Gemfile
- Gemfile.lock
- Rakefile
- README.md
- unicorn.rb

What we support:

- kube-gen files
- continuous integration files (with containers)
- documentation (markdown)
- databag secrets and keys
- development (legacy)
- AWS IAM roles
- internal tool configuration (UI/UX)
- legacy service discovery configuration
- project ownership and metadata
- .. and more!

A single deploy process for every change



@MELANIECEBULA

Deploy Deploy all changes Open a PR and merge your code to master atomically

GENERATING SERVICE BOILERPLATE



Reduce service boilerplate WITH GENERATORS!

- generators make best practices the default
- collection of config generators and language-specific generators
- uses thor for better control (update, review, commit)
- set up a "hello world" service with just one command

Reduce time to "hello world"



2 weeks

@MELANIECEBULA



2 minutes



- Store configuration alongside project code
- Abstract away your infrastructure with generators
- Abstract away complex tooling with a wrapper CLI
- Configuration and tools should set defaults
- Standardize on one process for storing and applying configuration changes

