Forced Evolution: Shopify's Journey to Kubernetes

Shopify

3000+
Employees

$26B
processed ‘17

600k+
merchants

80+k
Peak RPS

goto 2016

Running services.. everywhere

DCs	AWS PCI	AWS	Heroku
Chef+docker	Chef	Chef+???	

Service Tiers

More mature in SDLC

Greater business importance

Higher SLO

Earlier in SDLC

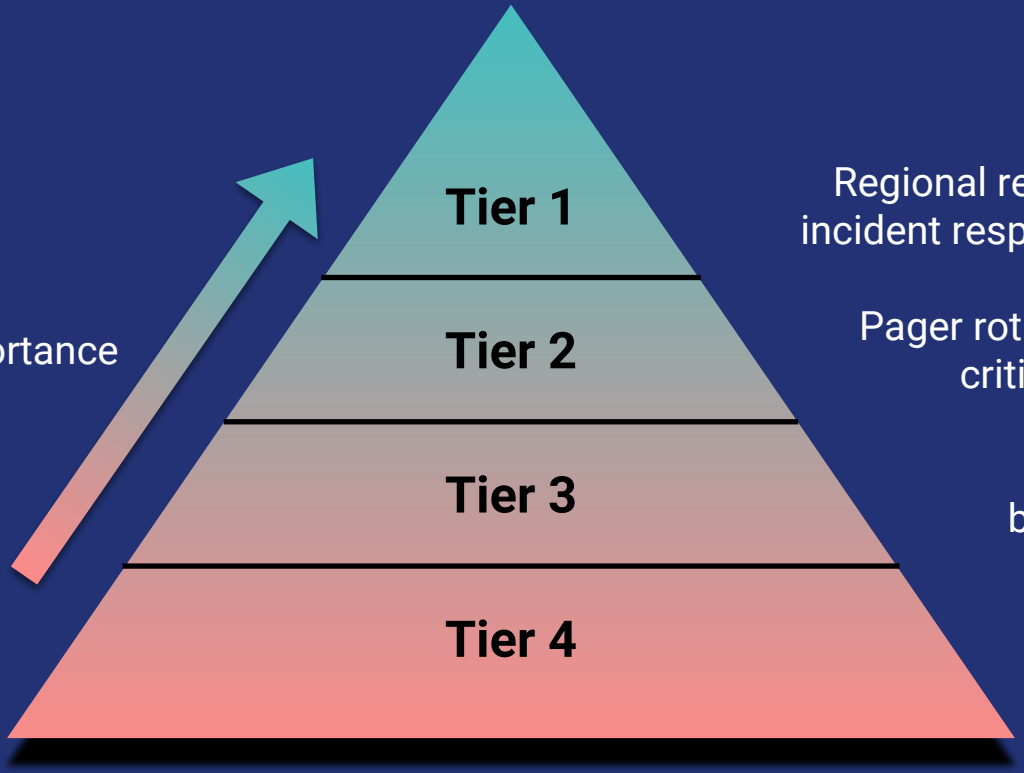

Regional redundancy, incident response drilling

Pager rotation, automated critical alerting

CI, Pingdom, backups, logging

Fewer requirements to encourage rapid prototyping

Not scalable

Things that won’t scale

- Manual / Artisanal processes
- Slow things/processes that make people wait
- Rusty knobs that don’t work when needed
- Wobbly things that don’t work first-time, every-time

Things that will scale

- Tested infrastructure
- Automation that works as expected, every time
- Give devs ability to self-serve with safety
- Train people to be experts in the systems they operate

Building a PaaS

“One Ring to rule them all, One Ring to find them, One Ring to bring them all and in the darkness bind them”

- The Lord of the Rings

Three principles

Paved road

Hide complexity


17超值
双层阿拉斯加鳕鱼堡
点击屏幕开始点餐
TOUCH TO START
轻松三步
送餐到桌
1.自助点餐
2.侧面取号
3.坐等美味
送餐到桌

Self serve

Why Kubernetes?

- Best traction of the open source projects
- Platform agnostic
- One of the most extendable solutions
- Written in Go
- Offered as a service in Google Cloud

Building blocks of running an application

- How to specify your apps runtime
- How to build your app
- How to deploy your app
- How to set up your dependencies

Creating application environment

Services DB

- Web UI for developers
- Application catalog
- Generation of Kubernetes manifests
- Configures builds and CI

Groundcontrol

- Go app living on clusters
- Creates k8s namespace
- Creates encryption keys
- Service accounts

cumulus-cat / niko-test

Manage expiry & deletion | Create another environment | Create Cloud Runtime | Environment 2

A CloudPlatform Test App :cloud: :cat:

Owners Edit
Product/Service line: Unclassified
@KnVerey — Director: Camilo Lopez

Runtime
Runtime not detected (Read more)

Production
Development
GitHub

Dependencies
Gem dependencies: 112
Missing licenses: 7

Service Graph (View)
No service tracing found in the last hour

---

shopify-services commented on 19 Apr

Owners: @KnVerey
Service: cumulus-cat/configdemo

Welcome to Shopify Cloud!

This PR contains the Kubernetes files and Shipit script you will need to deploy this app.

Please review this PR before merging it!

- [ ] Review the kubernetes templates to familiarize yourself with the kinds of configuration stored here
- [ ] Make sure `kube-shell.yml` contains everyone needing Kube Shell access to your containers
- [ ] We've configured the public key in this `config/secrets.<env>.ejson` on our cloud servers. If you have a previous public key in place, you will need to re-encrypt your secrets.
- [ ] If you need to add environment variable secrets (e.g. you are migrating from Heroku), you can add them to the "environment" key of `config/secrets.<env>.ejson`. Don't forget to run `ejson encrypt` before committing!

Questions about this PR? Read the docs or hop in the `#servicesdb` Slack channel. We're happy to help!

Stack

Rails (default)
Other

Classification

Your app is classified by default to be in the `tier4` cluster 🙌, but you can change the cluster by pressing the button below.

Advanced Options

Expiry Date

If your service is in `Tier4` or `Hackdays` cluster you need to set an expiry date.
Expired services will be candidates for shutdown and their owners will be notified before anything is removed.

22/0

Default is one month from today

Deployment Templates

Select the deployments that are needed for your application, note that each selection will be included in a pull request on your service's repo.
You'll have to review each pull request to complete your runtime setup.

For more information visit the **docs**

Workloads

- [ ] asset-uploading — So that you can have fast assets on a CDN. Rails apps need to add the 'shopify-cloud' gem to use this.
- [ ] db-migrate — So that you can run your migrations.
- [ ] jobs — So you can run background job workers
- [ ] scheduler-clock — Infrastructure needed to use [clockwork](https://github.com/Rykian/clockwork), a ruby gem for running tasks at a regular interval.
- [ ] web — This is so that you can actually run a web server.

Utilities

Buildkite + PIPA

- Buildkite acts as coordinator for Pipa
- Pipa agent builds Docker images
- Herokuish, Dockerfile, or custom build pipelines

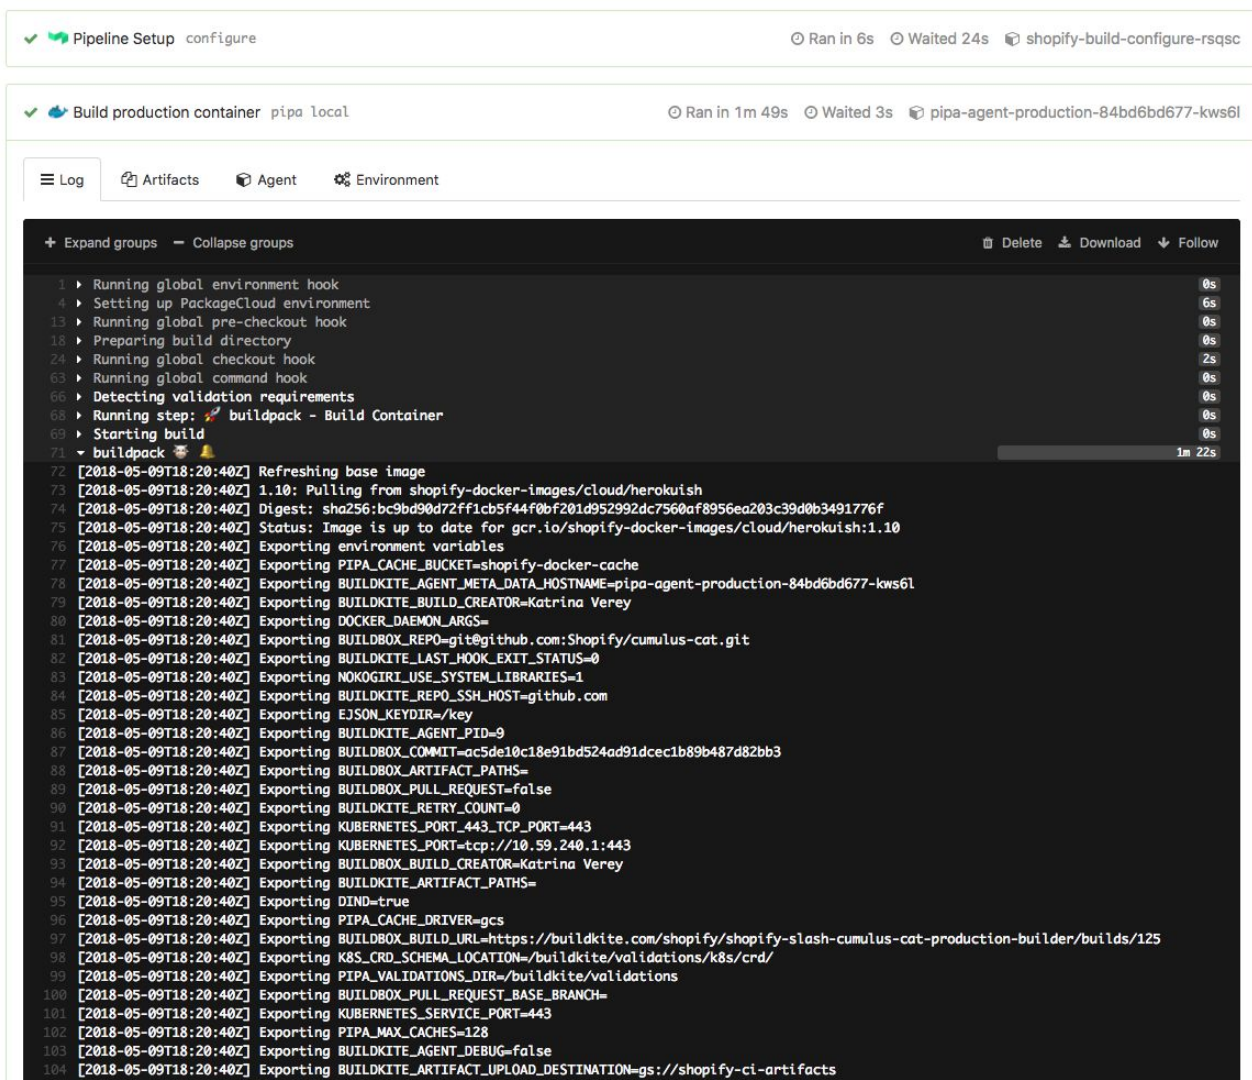

Builder Stats

6,000

average builds per weekday

450,000

images in GCR

kubernetes-deploy

- Pass/fail results on deploys
- Pre-deploy for ConfigMap/Secrets
- Protecting namespaces
- Pluggable

Cloudbuddies

- Create DNS records
- Fetch SSL certificates
- Create buckets, databases, services etc
- Set user editable quotas
- Set security rules
- Delete bad nodes

```
n1ko at n1pro in ~/Documents/shopify/go_workspace/src/github.com/Shopify/cloudbuddies on ±nodebuddy ✓
⬨ ls -lha|grep buddy                                                    ⎈ tier4/default
drwxr-xr-x   8 n1ko  staff   256B May 24 14:14 accountabilibuddy
drwxr-xr-x   9 n1ko  staff   288B May 24 14:14 bucketbuddy
drwxr-xr-x  18 n1ko  staff   576B Jun  8 09:10 cloudsqlbuddy
drwxr-xr-x   4 n1ko  staff   128B May 24 14:14 eventbuddy
drwxr-xr-x   6 n1ko  staff   192B May 24 14:14 identibuddy
drwxr-xr-x   6 n1ko  staff   192B May 24 14:14 kafka-access-buddy
drwxr-xr-x  10 n1ko  staff   320B May 24 14:14 mailbuddy
drwxr-xr-x  12 n1ko  staff   384B May 24 14:14 memcachedbuddy
drwxr-xr-x  11 n1ko  staff   352B Jun  8 09:10 namebuddy
drwxr-xr-x  10 n1ko  staff   320B May 30 10:48 netpolbuddy
drwxr-xr-x   7 n1ko  staff   224B Jun 12 16:44 nodebuddy
drwxr-xr-x   6 n1ko  staff   192B May 24 14:14 pingdomshopbuddy
drwxr-xr-x  10 n1ko  staff   320B Jun  8 09:10 rbacbuddy
drwxr-xr-x  14 n1ko  staff   448B Jun  8 09:10 redisbuddy
drwxr-xr-x   6 n1ko  staff   192B May 24 14:14 resourcebuddy
drwxr-xr-x   9 n1ko  staff   288B May 30 15:15 searchbuddy
drwxr-xr-x   8 n1ko  staff   256B May 24 14:14 secretbuddy
drwxr-xr-x   9 n1ko  staff   288B May 24 14:14 somebuddy
drwxr-xr-x   8 n1ko  staff   256B Feb  7 10:05 statefulbuddy
drwxr-xr-x  12 n1ko  staff   384B Jun  8 09:10 topicbuddy

n1ko at n1pro in ~/Documents/shopify/go_workspace/src/github.com/Shopify/cloudbuddies on ±nodebuddy ✓
⬨ ls -lha|grep buddy|wc -l                                              ⎈ tier4/default
      20
```

Namebuddy flow

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: web-nginx
  annotations:
    kubernetes.io/ingress.class: nginx
    kubernetes.io/tls-acme: 'true'
  labels:
   name: web
    app: cumulus-cat
    env: production
spec:
  tls:
  - hosts:
    - cumulus-cat.shopifycloud.com
....
```

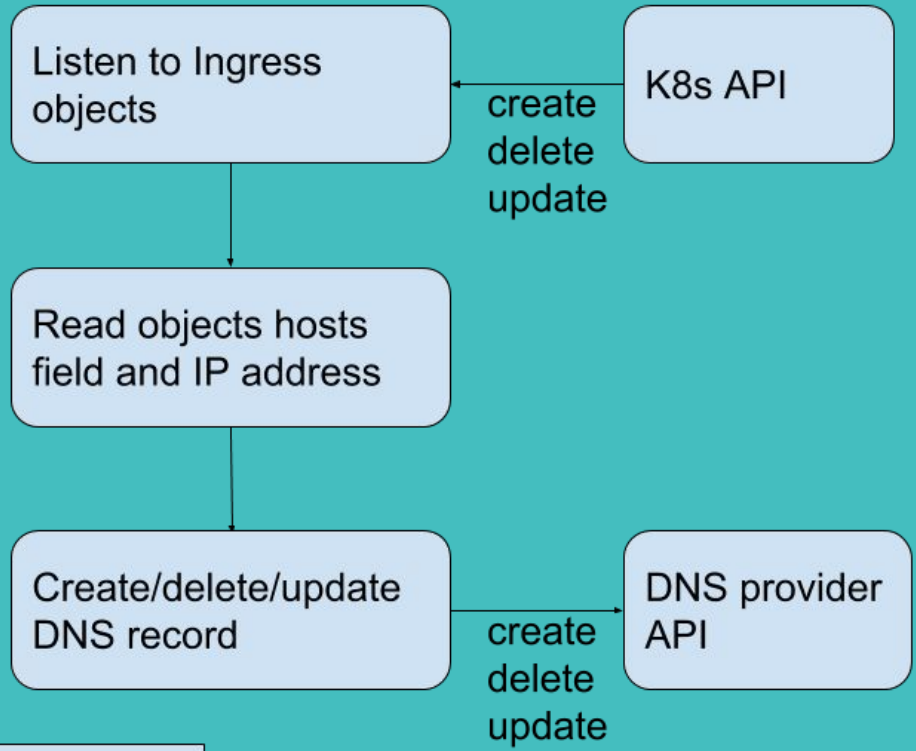


```
;docs.shopifycloud.com.             IN      A

;; ANSWER SECTION:
docs.shopifycloud.com.      59      IN      CNAME
cloud-docs-production-nginx-tier3.shopifycloud.com.
cloud-docs-production-nginx-tier3.shopifycloud.com. 59 IN A
35.185.75.173
```

CumulusCat

Niko

Secure | https://cumulus-cat.shopifycloud.com/#

Apps graphs cloud PR qcon registry Pull requests internal Okta Splunk Other Bookmarks

Home

Crash course to buddies

Extending k8s

- API's are well documented (if not super stable)
- Client libraries are high quality (at least on client-go)
- We can both extend functionality of current concepts (deployments, endpoints etc) but also create our own (CRDs)
- Distributed systems primitives (leader election, latches ...)
- These apps are be pure Go so they are unit testable, running and deployed as normal apps etc.

Kubernetes Controllers

An active state reconciliation process

- Watch desired and current state
- Try to mutate desired to current

```
for {
    desired := getDesiredState()
    current := getCurrentState()
    if desired != current {
        reconc(desired, current);
    }
}
```

Writing a controller

Workflow is always the same

- Authenticate to the cluster
- Create a watcher for events of specified type
- Implement functions to handle ADD/DELETE/UPDATE
- Profit!

Custom Resource Definitions

- Extend native k8s objects with your own abstractions
  - Eg. Memcache, Redis, Mail, MyFancyThingy
- Used by your own controllers to consume configuration params and doing something based on it
- Just like normal k8s resources like Deployment or Service

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment
  labels:
    app: nginx
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:1.7.9
        ports:
        - containerPort: 80
```

```
apiVersion: stable.shopify.io/v1
kind: Elasticsearch
metadata:
  name: <%= @app %>
  labels:
    app: <%= @app %>
    environment: <%= @env %>
    component: elasticsearch
spec:
  elastic-search-version: '6'
  zones:
    - us-east1-b
    - us-east1-c
    - us-east1-d
.....
```

```
......
  elasticsearch-spec: |-
    reindex.remote.whitelist: 10.*.*.*:9200
  node-specs:
    - replicas: 3
      cpu-limit: "1"
      mem-limit: 2G
      data-volume-size: 10Gi
  snapshot:
    bucket-name: shopify-<%= @app %>-<%= @env[0..3]
%>-es-snapshots
```

Supporting users

Workloads
Managed
Ephemeral

Network
Ingresses
Services

Configuration
ConfigMaps

Storage
Buckets

Datastores
Cloudsqls
Elasticsearches
Memcacheds
Redises

Utilities
Mails

Managed workloads

All | Deployment

backgroundqueue-high-priority	3/3 pods ready	Created 7 months ago
backgroundqueue-long-running	5/5 pods ready	Created 7 months ago
backgroundqueue-low-priority	25/25 pods ready	Created 7 months ago
cloudsql-07b3a605-b9b2-11e7-bb59-42010af00076	2/2 pods ready	Created 8 months ago
memcached-app-reviews-production	1/1 pods ready	Created 8 months ago
redis-e08649e1-b9b2-11e7-bb59-42010af00076	1/1 pods ready	Created 8 months ago
scheduler	1/1 pods ready	Created 7 months ago
vividcortex-07b3a605-b9b2-11e7-bb59-42010af00076	1/1 pods ready	Created 8 months ago
web	12/12 pods ready	Created 7 months ago

CloudPlatform guide

Search

- ***Getting started***
  - Help
  - Introduction
  - Prerequisites
  - Rails walkthrough
  - Static app walkthrough
  - Other apps walkthrough
- ***Troubleshooting***
  - Understanding Deploys
  - Scaling your app
  - Viewing Logs
  - Viewing Metrics
  - Additional Help
- ***Settings***
  - Restart an app
  - Delete an app
  - Tier Migrations
- ***Workloads***
  - Job workers
  - Web
- ***Datastores***
  - CloudSQL
  - Elasticsearch
  - Kafka
  - Memcached
  - Redis
- ***Network***
  - Ingress
  - SSL
- ***Security***
  - Authorization
  - Networking
  - Permissions

Overview

The CloudPlatform runtime is a unified deployment stack for Shopify applications backed by Kubernetes that strives to provide an out-of-the-box experience when possible (see Project Brief).

To use this documentation:

- Browse the categories in the sidebar or enter key terms in the search bar at the top right.
- Read the introduction for a primer if you're just getting started.
- Check out the glossary for clarification on any terms. You may also mouse over terms with a hatched underline to see a quick description.
- If you have read through the docs and still need help (or have an emergency) check out the guidelines on how to get help

Note

This documentation is intended to help make CloudPlatform self-service, without your needing to know too much about how the platform itself works. For knowledge sharing, and to better understand the platform, however, details about how things work are often provided to demystify some of the magic.

Documentation

Niko Kurtti 12:59
oncall cloudhelp

spy APP 12:59
Currently on call for Cloud-Help:
Cloud help (primary): @jenna (jenna.black@shopify.com) Until: 2018-06-14T21:00:00Z UTC
Cloud help (secondary): @stefan.budeanu (stefan.budeanu@shopify.com) Until: 2018-06-14T21:00:00Z UTC

Important: spy did not page Cloud-Help, use `spy page Cloud-Help [message]` to page them.

Report card

"The turn around time to getting an app running on cloud platform is unreal, you folks have really nailed it."

Challenges for developers

- How does my builds/deploys/everything work?
- How do I scale ?
- How do I debug?
- Is this worth it?

Challenges for SREs

- Giving up control over underlying infrastructure
- Container-only world and new tooling
- Customising the one platform to fit all needs
- Constant pressure to migrate apps
- Learning

Takeaways for building your own PaaS

- Target hitting eg. 80% of use cases
- Create patterns and hide complexity (but don’t restrict)
- Educate
- Get people excited
- Be conscious of vendor lock in

Future

- Polishing our tooling
- Making sure our platform keeps scaling and stable
- Optimising cost
- Multi cloud
- Service mesh

Thanks!

shopify

- github.com/Shopify/kubernetes-deploy
- github.com/Shopify/kubeaudit
- github.com/Shopify/shipit-engine

- https://www.flickr.com/photos/tomronworldwide/23953051439
- https://www.flickr.com/photos/cogdog/15152251297
- https://www.flickr.com/photos/jeffeaton/6586676089
- https://www.flickr.com/photos/27718575@N07/2683640267/in/photolist-569mKM-84HbTK-dtazDZ-iirKLf-2TEJmK-568rcD-6nofuM-9vLLH3-mUwPUR-9WhPqM-aqYH23-4JjwJx-6yLyB6-eaSpAu-nA38Vf-dCbp2o-56b387-8ekDpj-TEvNAr-op7reD-THmXQN-SBT2KU-QHezTj-SNuQzQ-c21rtC-pypWsn-fFRbW3-6YJuy4-fLWsf7-56dt27-56cnzW-7oYTG6-bUA74H-a9cDgi-9SGPxs-5fGdyo-7VRDXn-GiGAKB-568Z9H-5FVvF7-oD2WF-8KyzR9-avherm-4KXUjb-e8XabH-nVMfaF-569fXV-h11V7-rByx-66uNnq
- https://commons.wikimedia.org/wiki/File:Self-service_kiosks_at_McDonald%27s_Cuiwei_Store_(20170427201418).jpg
- https://commons.wikimedia.org/wiki/File:Building_foundation.jpg
- https://commons.wikimedia.org/wiki/File:Pacific,_WA_%E2%80%94_New_house_under_construction_%E2%80%94_02.jpg