

MIGRATING INTERNET ARCHIVE TO KUBERNETES

KUBECON 2018 - SEATTLE, WA

@tracey_pooh
@dvanduzer

2018 December 13

 for keyboard shortcuts


<https://archive.org/~tracey/slides>



* Internet

Kubernetes
or Bust!

CHALLENGE OF BIG CHANGES



WHAT'S IN IT FOR US?

- language / platform / OS flexibility of Docker
- model to enforce rebuildable recipes
- more coding; less system building and admin
- avoid cpu excess/starvation from static resource allocation
- fewer "pet" servers



INTERNET ARCHIVE

A DIGITAL LIBRARY

CURATE

PRESERVE

CIRCULATE

CURATE

CURATE ALL THE THINGS!

- 300B+ web pages [WayBack Machine](#)
- 15M books
- 4M videos & TV News programs
- 4M audio & concerts
- 3M images
- 200K software items & emulation

Find It

Product Information

Customer Support

Technology & Research

Developer World

Groups & Interests

Resources Online

About Apple

Apple Sites Worldwide

- Switzerland
- Taiwan
- Turkey
- UK & Ireland
- United States

Go

Where to Buy

Register to Win

Software Updates

Home Page Archives



Welcome to Apple



Introducing CyberDrive

Register today for a free CD-ROM.

EMATE 300

Mobile, Affordable, & Smart



MOVIES FROM MARS



QuickTime VR Takes You Out of this World

What's Hot

Preorder Mac OS 8

Now you can [preorder Mac OS 8](#), described by Macworld as "the most comprehensive update to the Mac OS in years, sporting a bold new look, a speedier Finder, more shortcuts and integrated Internet functions."

Want a PowerBook?

Qualify to win a [PowerBook 3400/200](#) by [entering](#) this month's Apple Registration Sweepstakes.

Big Help for Small Biz Find out what you need to make your small business a success. [Apple Small Business](#) has information about everything from great deals on leasing to

Be the First to Know


Learn about new Macintosh software releases the moment they become available. Check [Hot Mac Products](#) to hear about programs like Speed Demon, ReBirth RB-338 and QuickCRC.

Newton Connects

Newton, Inc., will enhance network connectivity for Newton-based devices this fall via [Newton Internet Enabler 2.0](#). Ethernet capability can connect devices to Local Area Networks.

Welcome to Mars

See alien terrain in [QuickTime VR movies from](#)




Goody Two-Shoes

 Edit  Manage  History  Book view  Barcode

Publication date c1888

Topics [Brothers and sisters](#), [Orphans](#), [Conduct of life](#), [Education](#)

Publisher New York: Melrose Publishing Co.



Pac-Man

 Edit  Manage  History  Barcode

by [Namco Limited](#)

Publication date [1983](#)

Secure https://archive.org/details/Sita_...

Search



1	DVD ISO (Movie only)	01:21:31
2	1080p HD compressed	01:21:50
3	1080p HD uncompressed	01:21:50
4	Sita Sings the Blues	01:21:31



Sita Sings the Blues

 Edit  Manage  History  Barcode

by [Nina Paley](#)

Publication date [2008](#)

Usage [CC0 1.0 Universal](#)  

Source File Authoring File can be found here: https://www.archive.org/details/Sita_Sings_the_Blues...

LOW FRICTION MIGRATIONS TO KUBERNETES

- book scanning apps
- crawl frontier management
- finite duration projects
- 'green field' / standalone apps

PRESERVE

PRESERVE

- 50+ petabytes over 22 years
- Fourth generation storage system
- Independence by self-hosting
- Multi-datacenter replication



SIMPLICITY == LONGEVITY

- simple directories of files
- original files with metadata
- multiple checksums
- changes logged forever
- replication: 'just' rsync

```
<file name="commute.mp4" source="derivative">
  <title>commute time-lapse</title>
  <format>h.264</format>
  <original>commute.avi</original>
  <md5>ff17ed66e7db5693dd208dd6ac488ff8</md5>
  <mtime>1325973601</mtime>
  <size>11919082</size>
  <crc32>ad1df03a</crc32>
  <sha1>e9f9de8379cd25653d487ab30d198fc61a050091</sha1>
  <length>115.61</length>
  <height>480</height>
  <width>640</width>
</file>
```

- ~~striped RAID~~
- ~~distributed filesystems~~
- boring block devices

CIRCULATE

SERVING

- 300B+ web pages [WayBack Machine](#)
- 15M books
- 4M videos & TV News programs
- 4M audio & concerts
- 3M images
- 200K software items & emulation

...TO ~4M DAILY VISITORS

TIMELINE / EVOLUTION

Month	Progress
2014/11	docker for audio fingerprints - controversial
2015/10	MozFest OpenNews docker talk mind blown
2016/8	switched mp3 / mp4 generation to docker
2017/6	IA Engr. Presentation
2018/6	IA Engr. Special Topic talk
2018/7	GitLab 'Auto DevOps' Changes Everything - Dev & Ops Harmony - Confessions of a middle child

TIMELINE / EVOLUTION

Month	Moves into GitLab Auto DevOps Kubernetes
2018/8	archive.org main repo CI/ <code>test</code> migrated
2018/9	archive.org full pipelines, website 'review apps'
2018/10	bookscan webapps
2018/10	moved to GitLab's docker registry
2018/12	dweb.archive.org (archive.org decentralized)

LOGJAM BREAKTHROUGHS

- TeamUp! (end of 2017)
 - David - ops-heavy history 
 - Tracey - dev-heavy history 
- GitLab advertises **with Kubernetes** July/2018
 - **already using** on premise **GitLab**
 - w/ ssh-key user accounts .org-wide

'ON PREM' K8 CLUSTER FROM SCRATCH

- <https://gitlab.com/internetarchive/kre8>



- **kubeadm** to create cluster
- sets up GitLab 'Auto DevOps' CI/CD full pipeline
- need 1+ VM w/ ssh access and sudo
 - ubuntu based so far
- coworker can try - low committment
- **kallr**

AVOID BAD BLOOD



WORK TOWARDS COMMON GOALS

Devs & Ops



WWW-NAME (BESPOKE) -V-
REVIEW APPS (GITLAB + K8)

webapps quick development at archive.org

WWW-TRACEY.ARCHIVE.ORG

- ssh 'bastion' with homedirs
- NFS /home/ automounts throughout cluster
- edit:
 - emacs/vi on (linux) bastion
 - laptop editors sftp / rsync-over-ssh on [save]
- nginx virtual hosts
 - <https://archive.org> => /petabox/www/.. (production)
 - <https://www-NAME.archive.org> => /home/NAME/petabox/www/.. (dev)

WWW-NAME.ARCHIVE.ORG ISSUES

- single 'dest' (multiple branches tricky)
- ssh keys and access
- NFS /home/ dirs

K8 + AUTO DEVOPS 'REVIEW APPS'


- contractors can work/test full site
 - no ssh keys/access
 - 12h timezone shift OK
 - no staff/approval need - **game changer!**
- git branches => 'review apps' consistent/update over time
- **kre8/branch-sync-edit.sh**
 - 'copy file change to pod on save' fast script
 - webapp instant f/b w/o full push and CI/CD

AUTO DEVOPS PIPELINES

passed Pipeline #54 triggered 12 hours ago by  tracey

6 jobs from `master` in 16 minutes 57 seconds (queued for 2 seconds)

Pipeline Jobs 14



CRAWLING CASE STUDY

CRAWL DISCOVERY LOOP

1. Start with a list of websites
2. Fetch all those URLs
3. Extract Hyperlinks from all the fetched content
 - Now you have a list of websites again
 - Which you can Fetch and Extract and...

HERITRIX

- open source Java crawler developed at IA
- very good at pipelining these stages
- but, single machine operation

CRAWLHQ

- distributed frontier management
- multiple Heritrix instances share single URL deduplication queue

DEMO OF K8S AUTO-SCALING

REWRITING CRAWLHQ


- Kafka partitions for work queue
- FoundationDB for hashes of seen URLs
- 38,000 lines of code is now 73 lines of Python
- throughput scales automatically with GitLab **REPLICAS** variable
















**my god ..
it's full of.. petabytes**

CONTAINING A MONOLITH

- 4.8G tree => 6GB .. 9GB docker images
- `test` before `build` = faster dev feedback
- naive `git clone` every CI/CD phase = nonstarter

PETABOX REPO DOCKER IMAGES

[petabase]

Baseline docker image

Ubuntu LTS

pkgs

js npm/yarn

php pecl/pear

python pip2/3

petabox repo git **SHALLOW** clone

CI **test** after git fast fwd

[petabase]



[workspace]

[webnode]

[dev]

[homeno]

3rd party
deriver pkgs
+ config
ffmpeg ..

git fast fwd
archive.org
website
nginx+
supervisor

git fast fwd
empty DB/SE
daemons/supervisor
archive.org website
item store/process

git fast f
misc noc



[worker]

git fast fwd
mp3/mp4/etc

MONOLITH REPO SPEED++

- **[petabase]** baseline image recut every 1-4 weeks
 - git **shallow** clone (2/3 size)
 - all usages 'fast forward' to branch & commit
- **build** phase uses prior build
 - via registry docker pull for branch

HOW TO FAST FORWARD QUICKLY FROM WEEKS OLD GIT SHALLOW TREE?



* MONOLITH REPO GIT *

fast forward a git clone depth 1

- `git checkout --detach *`
- `git fetch --depth 1 -f -origin [brnch|ca6745]`
- `git checkout -f [brnch|ca6745]`

* avoid tangle with git tree state

- Workaround 'advertised' commits issue (self-hosted GitLab)
 - `omnibus gitconfig['system'] = {"uploadpack" => ["allowReachableSHA1InWant = true"]}`
 - then able to fetch any commit

BECAUSE WE HAVE **THIS** GUY (AND YOU DON'T)



Mark Sullivan *Web Archiving Software Engineer*

Mark joined the Archive in August 2016 to help Arc of the Archive, with new data transfer tools. He graduated with a Bachelor's in Computer Science and a Bachelor's of Computer Education (classes!). He has worked in different enough environments that he occasionally dreams of expressive computers, and

^ git pro

TOP TIPS

- can use simple `localstorage` for 1 or 2-node k8 clusters
- `POSTGRES_ENABLED=false` CI/CD var
 - avoids extra (unused!) PV per branch/deploy
- GitLab API for last master `test` phase status
 - 'OK to push live?'
- periodically:
 - remove old replicaset
 - trim registry

LOADBALANCER VS INGRESS FOR ALT PORTS

- `kubectl expose -ndweb deployment/production -type=LoadBalancer --name=porter`
- `kubectl patch -ndweb svc/porter -p='{"spec":{"externalIPs":["'${K8_INGRESS_IP?}"]}}'`
- allows for non :80 :443 :5000 ports
 - websockets eg:
 - `wss://dweb.me:4246`

NEAR FUTURE

archive.org production website => K8

now static list of VMs + haproxy

new elastic demand-based scaling

THE INTERNET ARCHIVE

TO COLLECT & TO SERVE

THE END