



Software Engineering Conference Russia 2018

October 12-13
Moscow

On declarative configuration management using Kubernetes and Helm

Alexander Chistyakov

ITMO University

First things first, what is DevOps?

- I will make a short break for you so you can ship a release
- And if your process is solid enough, you can ship while I talk
- Or even better, a release can ship itself when it's ready without even bothering you
- Do you need a couple of minutes just to check if something is shipping right now?

BTW, who are you?

- Are you a human being?
- Maybe a DevOps engineer?
- Maybe a cyborg?
- Maybe even a Jenkins node?
- If so, how many jobs you can run in parallel?

What does a CM process look like?

- We define a desired state declaratively

What does a CM process look like?

- We define a desired state declaratively
- We instruct a CM master tool to get a system to that state idempotently

What does a CM process look like?

- We define a desired state declaratively
- We instruct a CM master tool to get a system to that state idempotently
- Some CM magic begins

What does a CM process look like?

- We define a desired state declaratively
- We instruct a CM master tool to get a system to that state idempotently
- Some CM magic begins
- Midnight suddenly comes

What does a CM process look like?

- We define a desired state declaratively
- We instruct a CM master tool to get a system to that state idempotently
- Some CM magic begins
- Midnight suddenly comes
- Infrastructure gets to the desired state!

What does a CM process look like?

- We define a desired state declaratively
- We instruct a CM master tool to get a system to that state idempotently
- Some CM magic begins
- Midnight suddenly comes
- Infrastructure gets to the desired state!
- Sounds not too difficult, should be quite easy to implement

A bit of history

- Are you familiar with Puppet CM system?
- Raise your hands if you like it! (But how on Earth is it possible?)

A bit of history

- Are you familiar with Chef CM system?
- Raise your hands if you like it! (Please don't cry, SECR is a safe place and I'm your friend)

A bit of history

- Are you familiar with Ansible CM system?
- Raise your hands if you like it!

If you know Ansible, pls help me a bit

- I need to configure a cluster of something
- No, not a cluster yet, just a monitoring tool for a cluster

Okay let's try

alexclear / ansible-graphite-playbook

Unwatch 1 Star 7 Fork 2

Code Issues 0 Pull requests 0 Projects 0 Wiki Insights Settings

An Ansible playbook to set up StatsD and Graphite Edit

[Manage topics](#)

2 commits 1 branch 0 releases 1 contributor MIT

Branch: master New pull request Create new file Upload files Find file Clone or download

alexclear First public version Latest commit e577c42 on 11 Apr 2014

graphite-supervisor	First public version	5 years ago
graphite	First public version	5 years ago
LICENSE	Initial commit	5 years ago
README.md	Initial commit	5 years ago


What's inside




```
107
108   - name: install make
109     apt: name=make state=present
110
111   - name: install cpanm
112     apt: name=cpanminus state=present
113
114   - name: install Net::Statsd::Server
115     cpanm: name=Net::Statsd::Server
116
117   - name: add a statsd user
118     user: name=statsd state=present
119
120   - name: put Net::Statsd::Server config
121     template: src=graphite/localConfig.js dest=/home/statsd/localConfig.js owner=statsd group=statsd mode=064
122
123   - name: put statsd service descriptor
124     template: src=graphite-supervisor/statsd.conf dest=/etc/supervisor/conf.d/statsd.conf owner=root group=rc
125     notify:
126       - restart supervisor
127
128   handlers:
129     - name: restart supervisor
130       service: name=supervisor state=restarted
```

What's wrong with it?

- Nothing, this module was just awesome
- Except that it was poorly written, not documented at all, hard to use, not supported and finally abandoned
- Still 7 stars on GitHub!

Let's try again





 **kofonfor / ansible-role-graphite**
forked from [gitinsky/ansible-role-graphite](#)

 Unwatch **1**  Star **0**  Fork **1**

[Code](#) [Pull requests 0](#) [Projects 0](#) [Wiki](#) [Insights](#) [Settings](#)


No description, website, or topics provided. [Edit](#)





[Manage topics](#)

 **38** commits  **1** branch  **0** releases  **3** contributors

Branch: **master** [New pull request](#) [Create new file](#) [Upload files](#) [Find file](#) [Clone or download](#)

This branch is even with gitinsky:master. [Pull request](#) [Compare](#)


 **alexclear** Merge github.com:kofonfor/ansible-role-graphite Latest commit e6c38d9 on 12 Feb

 api-cyanite	Replace sudo with become	2 years ago
 api	Follow a symlink	9 months ago
 carbon	Merge github.com:kofonfor/ansible-role-graphite	8 months ago
 web	Extract graphite webapp_dir as a configuration parameter	2 years ago

What's inside


```
22 - name: put carbon service descriptor
23   template: src=carbon-eye.rb dest=/etc/eye/carbon.rb owner=root group=root mode=0644
24   when: "not ((ansible_distribution == 'Debian' and (ansible_distribution_version|version_compare(8.0, '>='))) or (ansible_dis
25
26 - name: create {{ graphite_prefix }}/storage/
27   file: dest={{ graphite_prefix }}/storage/ state=directory
28
29 - name: load config in eye
30   command: bash -c "source /usr/local/rvm/scripts/rvm && eye l /etc/eye/apps.eye"
31   become: yes
32   when: "not ((ansible_distribution == 'Debian' and (ansible_distribution_version|version_compare(8.0, '>='))) or (ansible_dis
33
34 - name: put systemd unit file to its destination
35   template: src=carbon.service dest=/etc/systemd/system/graphite-carbon.service owner=root group=root mode=0644
36   when: "(ansible_distribution == 'Debian' and (ansible_distribution_version|version_compare(8.0, '>='))) or (ansible_distribu
37   notify: systemctl daemon-reload
38
39 - name: enable graphite-carbon service
40   service: name=graphite-carbon enabled=yes state=started
41   when: "(ansible_distribution == 'Debian' and (ansible_distribution_version|version_compare(8.0, '>='))) or (ansible_distribu
```

What's wrong with it?



What's terribly wrong with it?

- Can you tell for sure what is a status of the Carbon systemd service?
- And if you can, imagine a situation where you provide a service unit file and then use Ansible “notify” mechanism to call a handler which in turn enables the service



What's terribly wrong with it?

- It's hard to explain (not so hard to implement though), but people tend to be extremely proud implementing solutions like this
- Why on Earth does Ansible allow this?
- **Please note that Ansible handlers run at the end of script execution**



BETTER SAFE
THAN SORRY

Back to the CM process for a moment

- We want to run a set of microservices on top of a monolithic Linux kernel
- If this still makes sense, we can consider services as building blocks

A dictionary for people from late 90s

- A microservice = a process
- A Docker image = a package
- Kubernetes = an operating system for microservices
- Helm = a package managing system (think apt in Debian or yum in RHEL) for Kubernetes

Let's try



Alex Chistyakov

alexclear

I opened Vim in 1997 and still did not find a way to quit it.

[Edit bio](#)

@kofonfor

Saint-Petersburg, Russia

alexclear@gmail.com

Organizations



Overview **Repositories 91** Stars 852 Followers 113 Following 61

k8s--lab

Type: All

Language: All

New

6 results for repositories matching k8s--lab

Clear filter

k8s-prometheus-grafana-lab

Prometheus/Grafana stack deployed to k8s

MIT License Updated just now

k8s-percona-xtradb-cluster-lab

Percona XtraDB Cluster in Kubernetes

★ 1 MIT License Updated 10 days ago

k8s-freehck-efk-lab

An EFK stack deployed to freehck's k8s installation

MIT License Updated 11 days ago



What's inside a lab

- A Vagrant-based Ansible-provisioned K8s setup
- 3 Vagrant boxes
- Using Kubespray (<https://github.com/kubernetes-incubator/kubespray>) to set up Kubernetes (Kubespray is where Ansible meets K8s)
- Using Helm to provision Kubernetes apps

What's inside a playbook

es (13 sloc) | 503 Bytes

Raw

Blame

History

```
- hosts: all
  become: yes
  roles:
  - role: local-persistent-volumes

- import_playbook: ./kubespray/cluster.yml

- hosts: kube1
  become: yes
  roles:
  - role: helm
  - { role: k8s-namespace, k8s_namespace: "stolon" }
  - { role: k8s-stolon-cluster, stolon_helm_release: "stolon", k8s_namespace: "stolon" }
  - { role: postgresql-grafana }
  - { role: k8s-grafana, k8s_namespace: "stolon", grafana_helm_release: "grafana", k8s_grafana_admin_name: "test",
```

Ansible calls Helm

s (4 sloc) | 312 Bytes

Raw

Blame

History

```
- name: put values
  template: src=values.yml dest=/tmp/stolon-{{ stolon_helm_release }}.yml

- name: upgrade {{ stolon_helm_release }} release using a Helm chart
  shell: helm upgrade {{ stolon_helm_release }} stable/stolon --install --namespace {{ k8s_namespace }} -f /tmp/sto
```

Tired of dictionaries (the roads must roll!)

- Ansible = Helm
- Ansible role = Helm chart
- Python = Golang (this goes far beyond the scope of this session)
- Jinja templates = some custom templating engine
- ____ = tiller

How Helm works

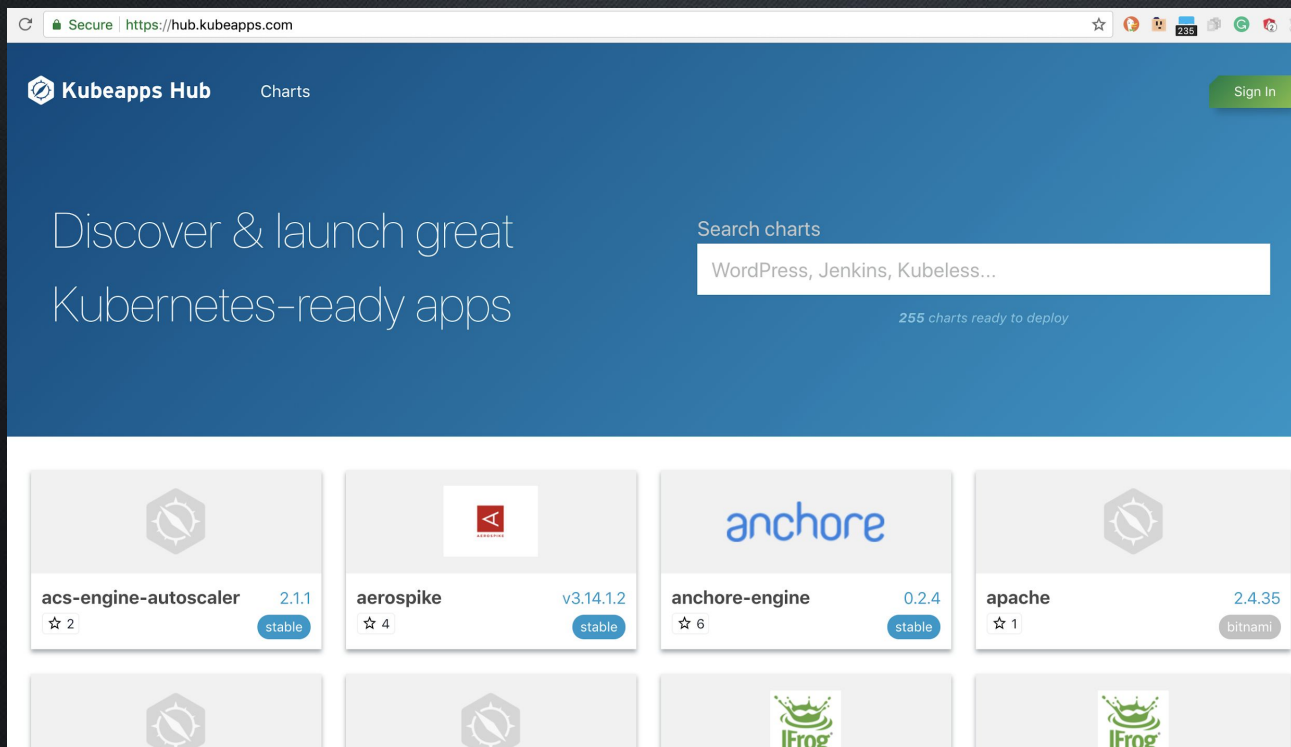
- Helm generated K8s YAML configs
- Helm transfers generated config files to the K8s cluster
- Tiller is an agent which runs in the cluster, effectively a cluster-side Helm service
- Tiller applies configuration files and labels a release
- Just like a Chef client 7 years ago but without a memory leak :)

Helm charts repository

- It's centralised
- It's on GitHub (<https://github.com/helm/charts>)
- It grows very fast
- If you submit a patch it will be landed in about a month, which I consider quite fast

Helm charts repository

- <https://hub.kubeapps.com> (a nice landing page for non-kinesthetic people)















The screenshot shows the Kubeapps Hub website. The header includes the Kubeapps Hub logo, the word "Charts", and a "Sign In" button. The main content area features the text "Discover & launch great Kubernetes-ready apps" on the left and a search bar on the right. The search bar contains the text "WordPress, Jenkins, Kubeless..." and a "255 charts ready to deploy" indicator below it. Below the search bar, there is a grid of chart cards. The first row contains four cards: "acs-engine-autoscaler" (version 2.1.1, stable, 2 stars), "aerospike" (version v3.14.1.2, stable, 4 stars), "anchore-engine" (version 0.2.4, stable, 6 stars), and "apache" (version 2.4.35, bitnami, 1 star). The second row shows the start of another row of cards, including one with the "JFrog" logo.

What's inside a chart

..		
📁 ci	[stable/stolon] Fix issue with secret deletion during helm upgrade (#...	4 months ago
📁 templates	[stable/stolon] Add option to force-disable stolon auto-cluster-creat...	2 months ago
📄 .helmignore	[stable/stolon] add stolon chart (#5490)	5 months ago
📄 Chart.yaml	[stable/stolon] Updated README (#7257)	a month ago
📄 README.md	[stable/stolon] Updated README (#7257)	a month ago
📄 values.yaml	[stable/stolon] Add option to force-disable stolon auto-cluster-creat...	2 months ago
📄 README.md		

Templates are just declarative configs

..		
 NOTES.txt	[stable/stolon] add stolon chart (#5490)	5 months ago
 _helpers.tpl	[stable/stolon] add stolon chart (#5490)	5 months ago
 create-cluster-job.yaml	[stable/stolon] Add option to force-disable stolon auto-cluster-creat...	2 months ago
 keeper-headless-service.yaml	[stable/stolon] add stolon chart (#5490)	5 months ago
 keeper-statefulset.yaml	[stable/stolon] Add support for priorityClasses (#6607)	3 months ago
 proxy-deployment.yaml	[stable/stolon] Add support for priorityClasses (#6607)	3 months ago
 proxy-service.yaml	[stable/stolon] add stolon chart (#5490)	5 months ago
 role.yaml	[stable/stolon] add stolon chart (#5490)	5 months ago
 rolebinding.yaml	[stable/stolon] add stolon chart (#5490)	5 months ago
 secret.yaml	[stable/stolon] Fix issue with secret deletion during helm upgrade (#...	4 months ago
 sentinel-deployment.yaml	[stable/stolon] Add support for priorityClasses (#6607)	3 months ago
 serviceaccount.yaml	[stable/stolon] add stolon chart (#5490)	5 months ago

What's wrong with Helm?

- Unfamiliar templating syntax (if you are not a Lisper)

```
lines (18 sloc) | 585 Bytes Raw
1  {{- if and .Values.alertmanager.enabled (empty .Values.alertmanager.configMapOverrideName) -}}
2  apiVersion: v1
3  kind: ConfigMap
4  metadata:
5    labels:
6      app: {{ template "prometheus.name" . }}
7      chart: {{ .Chart.Name }}-{{ .Chart.Version }}
8      component: "{{ .Values.alertmanager.name }}"
9      heritage: {{ .Release.Service }}
10     release: {{ .Release.Name }}
11     name: {{ template "prometheus.alertmanager.fullname" . }}
12  data:
```

What's wrong with K8s and Helm?

- They are written in Golang
- Really, seriously, please stop using Golang for everything!
- We are humans not rodents, we already invented nuclear bombing, Haskell and even soccer, why Golang after all???

Conclusions

- Classical CM systems were never good
- Kubernetes and Helm are not good too, but we are just not fully aware yet
- And, frankly speaking, I don't care



Contacts

- Alex Chistyakov
- alexclear@gmail.com
- +7-921-643-93-27 (Please don't call me, I am a sociophobe)
- @demeliorator (in Telegram and Skype)

