Flexible logging pipelines with Fluentd and Kubernetes

Jakob Karalus, IT-Consultant, codecentric

Log forwarding - the default way

- A daemonset for fluentd on every node
- Kubernetes writes logs to disk
- Fluentd read logs from disk
- Filters & enrich with metadata
- Routes to output

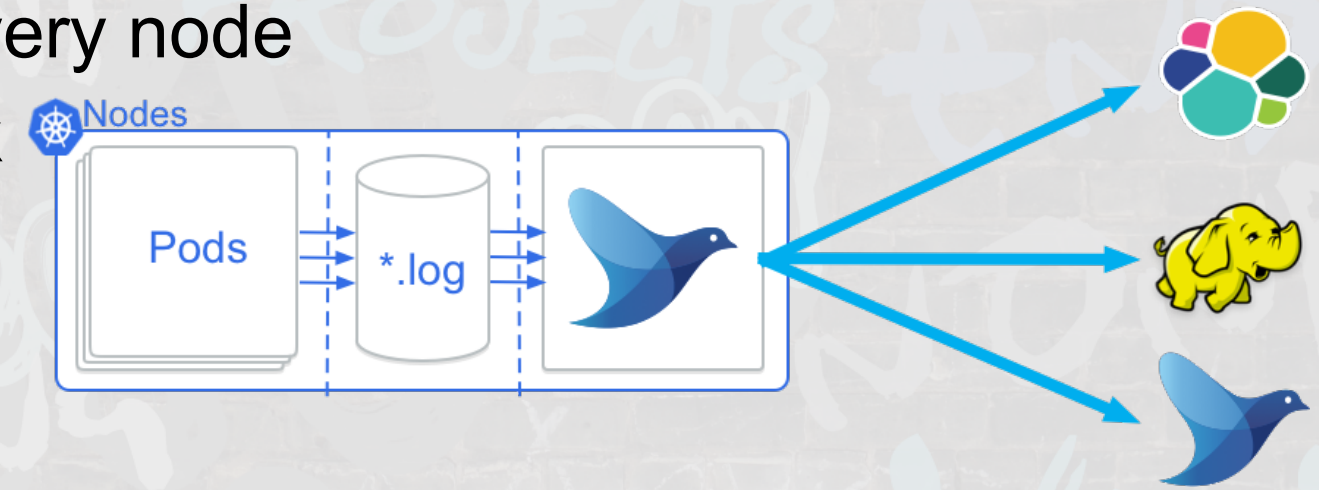

Problems

- Shared config for all teams -> coordination problem
- Hard to redeploy of daemonset for every change
- Quickly gets complicated for large shared clusters

Solutions:

- Solution A: Setup central Parsing/Routing “Point”
  - Coordination problems, extra infrastructure but better than static
- Solution B: Sidecars for every different Route/Parse
  - Duplicated resources/manpower but really flexible
- Solution C: Dynamic config generation from user defined attributes
  - Complex setup but flexible interface for devs

Dynamic Generation

1. Define config in Pod annotations
2. Create a templated fluentd config
3. Read annotations from API
4. Generate config with template
5. Reload fluentd config

```
apiVersion: v1
kind: Pod
metadata:
  name: busybox1
  annotations:
    fluentd_target: >
      [
          {
              "output": "elasticsearch",
              "host": "localhost"
          }
      ]
spec:
```

Create config using golang templates

```
{{ $target_pods := whereExist .Pods "ObjectMeta.Annotations.fluentd_target" -}}
{{- range $pod := $target_pods -}}
  {{/* get dockerid for the logfile */}}
  {{ $cID := trimPrefix (index $pod.Status.ContainerStatuses 0).ContainerID "docker://"}}
  {{/* parse annotation to readable config */}}
  {{ $config := first (parseJson $pod.ObjectMeta.Annotations.fluentd_target) }}
<match kubernetes.var.log.containers.{{ $cID }}.log>
   type {{ $config.output}}
   host {{ $config.host}}
</match>
{{- end }} #end pod for loop
```

Read API and render template

- Read annotations from API and create fluentd config:

```
fluentd  master±  kube-gen -watch -type pods -wait 2s:3s -post-cmd 'td-agent reload' td-agent.conf.tmpl rendered.conf
 | min wait time reached
 | refreshing state...
 | done. took 6.47362ms
 | running command [td-agent reload]
```

- Reload fluentd

Output

```
<source>
  type tail
  path /var/log/containers/*.log
  tag kubernetes.*
  format json
</source>

<filter kubernetes.var.log.containers.*.log>
  type kubernetes_metadata
</filter>

<match kubernetes.var.log.containers.36a533{SNIP}.log>
   type elasticsearch
   host localhost
</match>
```

End

- Small example: github.com/krallistic/kubernetes-fluentd
- Also done with Beats: github.com/kylemcc/kube-filebeat
- Better generator: github.com/kylemcc/kube-gen
- @krallistic, jakob.karalus@codecentric.de