MUX Treating Dashboards Like Code Scott Kidder, Staff Software Engineer @ Mux Grafanacon, February 26, 2019





- Background on Mux
- Monitoring for Mux Data
- Greenfield Monitoring Opportunities with Mux Video
- Goals for monitoring
- Questions





Background on Mux







What is Mux?

- Mux Data: Analytics for Video (2016)
- Mux Video: API for Video (2018)
- Mux Video makes it easy to publish video with a REST API call
- Optimal video encoding settings chosen automatically
- Deployments in AWS and Google Cloud





Monitoring for Mux Data





For a Moment, Let's Return to a Simpler Time

- Mux has used Grafana since inception (early 2016)
- Single deployment of Rancher container orchestration system in AWS
- Supported Mux Data, our only product at the time
- Single Grafana instance for all dashboards
- Single InfluxDB instance for application metrics













But in many ways, things were more difficult...





Problems began to surface

- Management of alerting rules was performed in Chronograf
- Ran a second visualization tool just to administer alerting rules
- No versioned history of alerting rules
- Rules were often disabled during a deploy or maintenance, and then people would forget to re-enable them, leading to undetected incidents
- Unclear why alerts were disabled, and whether it's safe to delete





Greenfield Monitoring Opportunities with Mux Video





Mux Video Development

- Late 2017 we began developing Mux Video
- with Mux Data
- Decided to run all services in Kubernetes and monitor with **Prometheus and Grafana**

We had already run some proof-of-concept Kubernetes cluster





GCE US East 1



GCE US East 4







Goals for Monitoring







Goals

- 1. Easily configure which services are scraped by Prometheus
- 2. Run policy checks on alert rules with each build
- 3. Store the dashboards and alert rules alongside code
- 4. Automatically deploy dashboards and alert rules to Kubernetes clusters each time we ship code





Goal #1 Easily configure which services are scraped by Prometheus







Prometheus Monitoring in Kubernetes

- Using the Prometheus Operator to configure Prometheus and Alertmanager
- <u>https://github.com/coreos/prometheus-operator</u>
- Uses Kubernetes label metadata to target which services to scrape and on which port





Prometheus: Kubernetes Service Monitor

1) Examine services in all Kubernetes names

2) Match on services with a "monitoring: core

3) Scrape whatever port is named "metrics"

	anillani an
	apiversion:
	monitoring.coreos.com/vl
	kind: ServiceMonitor
	metadata:
	name: core-servers
	namespace: monitoring
	labels:
	k8s-app: core-servers
	spec:
	jobLabel: core-servers
spaces	namespaceSelector:
•	any: true
	selector:
e" label	matchLabels:
	monitoring: core
	endpoints:
	- port: metrics
	interval: 10s
	honorLabels: true



Prometheus: Monitored Service

1) Simply add the "monitoring: core" label to a server







Services Scraped

Prometheus Alerts Gra	ph S	Status - Help		
Targets				
accesslogs-kafka (3/3	up)			
Endpoint	State	Labels	Last Scrape	Err
http://100.96.55.84:5556/metrics	UP	endpoint="metrics" instance="100.96.55.84:5556" namespace="gce-us-east1-production" pod="accesslogs-kafka-broker-0" service="accesslogs-kafka"	3.3s ago	
http://100.96.58.20:5556/metrics	UP	endpoint="metrics" instance="100.96.58.20:5556" namespace="gce-us-east1-production" pod="accesslogs-kafka-broker-1" service="accesslogs-kafka"	4.169s ago	
http://100.96.59.22:5556/metrics	UP	endpoint="metrics" instance="100.96.59.22:5556" namespace="gce-us-east1-production" pod="accesslogs-kafka-broker-2" service="accesslogs-kafka" service="accesslogs-kafka"	9.769s ago	

apiserver (3/3 up)

Endpoint	State	Labels	Last Scrape	Error
https://10.142.0.10:443/metrics	UP	endpoint="https" instance="10.142.0.10:443" namespace="default" service="kubernetes"	8.029s ago	
https://10.142.0.3:443/metrics	UP	endpoint="https" instance="10.142.0.3:443" namespace="default" service="kubernetes"	10.574s ago	
https://10.142.0.9:443/metrics	UP	endpoint="https" instance="10.142.0.9:443" namespace="default" service="kubernetes"	2.951s ago	







Goal #2 Run policy checks on alert rules with each build





Prometheus: Automated Policy Check

🗸 🙏 P	olicy Check docker-compose -f .buil
≡ Log	Artifacts D Timeline
+ Exp	oand groups – Collapse groups
1 ► 17 -	Preparing build folder
17 ×	docker-compose -f .buildkite/docker
19 C 20	hecking ./core/golang/kafka/monitor SUCCESS: 1 rules found
21 22 C	hecking ./core/servers/consul/monit
23 24	SUCCESS: 4 rules found
25 C	hecking ./core/servers/hadoop/monit
20	SUCCESS. S Fulles Tound
28 C 29	hecking ./core/servers/zookeeper/mo SUCCESS: 3 rules found
30	'heelving (dete (internel (servers (sh
32	SUCCESS: 7 rules found





Prometheus: Automated Policy Check

1) Use promtool to validate alert rules files

2) Verify that all files end with a new-line to allow for concatenation

```
"*.rules.yaml" | sort)
  rc=$?
YAML file."
    echo ""
    ERRORS="yes"
  fi
    echo
         ** **
    ERRORS="yes"
  fi
done
```

MONITORING YAML FILES=\$(find \$searchdir -type f -name for f in \$MONITORING YAML FILES; do promtool check rules \$f

if [[\$rc != 0]]; then echo "\$f is not a valid Prometheus alert rule

| tail -c 1) LAST CHAR= $(cat \fint \ | \ tr \ | \ n' \ | \ \#'$ if [[\$LAST CHAR != "#"]]; then echo "\$f does not end with a new line."



Goal #3 Store the dashboards and alert rules alongside code





Code Organization

1) Dashboards are named "*dashboard.json", and stored in a "monitoring/grafana" directory for the associated component

2) Alert rules are named "*.rules.yaml" and kept in a "monitoring" directory

✓ servers	
▶ autoscale	
✓ consul	
monitoring	
<pre>{} consul-dashboard.json</pre>	
consul.rules.yaml	
≡ consul.watches	
 README.md 	
! run-aws-us-east-1-production.yaml	
! run-aws-us-east-1-snowflake-staging.yaml	
! run-aws-us-east-1-staging.yaml	
! run-gce-europe-west1-production.yaml	
! run-gce-us-east1-production.yaml	
<pre>! run-gce-us-east4-production.yaml</pre>	
! run-gce-us-west1-staging.yaml	
! run-local.yaml	
! run.yaml	
▲ hadoop	
monitoring	
✓ grafana	
<pre>{} hadoop-hdfs-dashboard.json</pre>	
! hadoop.rules.yaml	
<pre>! run-aws-us-east-1-staging.yaml</pre>	
<pre>! run-gce-us-east1-production.yaml</pre>	
<pre>! run-gce-us-west1-staging.yaml</pre>	

run.yaml

MUX



Goal #4 Automatically deploy dashboards and alert rules to Kubernetes clusters each time we ship code





Automatic Deployment of Dashboards and Alert Rules

- for servers across all target environments
- and Prometheus alert rules
- Buildkite deploy plan applies Kubernetes manifests and ConfigMaps to each Kubernetes cluster
- Grafana and Prometheus ConfigMaps automatically reloaded

Our Buildkite builds automatically generate Kubernetes manifest

Also generate Kubernetes ConfigMaps with Grafana dashboards









done

- -type f -name "*.rules.yaml") for file in \$RULES FILES; do cp \$file \$RULES DIR
- for searchdir in "\${@:2}"; do RULES FILES=\$(find \$searchdir
- mkdir -p \$OUTPUT DIR rm -r \$OUTPUT DIR/* || true
- RULES DIR=\$1
- #!/bin/bash





1) Begin rendering a Kubernetes ConfigMap manifest

2) Concatenate contents of each alert rule file to the ConfigMap

Generate Kubernetes ConfigMap with Alert Rules

```
set -e
NAMESPACE=$1
OUTPUT FILE=$2
RULES DIR=$3
mkdir -p $(dirname $OUTPUT FILE)
cat <<-EOF > $OUTPUT FILE
apiVersion: v1
kind: ConfigMap
metadata:
  name: prometheus-k8s-rules
  namespace: $NAMESPACE
  labels:
    role: prometheus-rulefiles
    prometheus: k8s
data:
EOF
for f in $(find $RULES DIR -type f -name
"*.rules.yaml")
do
          $(basename $f): |+" >> $OUTPUT FILE
  echo "
  cat $f | sed "s/^/
                        /g" >> $OUTPUT FILE
done
```







- Prometheus Operator includes a config reloader that monitors the ConfigMap for changes
- Sends web hook to Prometheus instructing it to reload its config

Automatic Deployment of Prometheus Alert Rules





Gather Grafana Da	ishbo
	#!/bin set -e
	OUTPUI mkdir rm -r
I) Find all Grafana dashboard and datasource files conforming to naming pattern	for se f -nam datasc
	done
30	



oards and Datasources

n/bash

```
C DIR=$1
-p $OUTPUT DIR
$OUTPUT DIR/* || true
```

```
earchdir in "${@:2}"; do
DASHBOARD FILES=$(find $searchdir -type)
ne "*-dashboard.json" -o -name "*-
ource.json" | sort)
for file in $DASHBOARD FILES; do
  echo "FILE: $file"
      cp $file $OUTPUT DIR
  done
```



Render ConfigMap with Grafana Dashboards

Have been using the `grafana-dashboards-configmap-generator` script at https://github.com/eedugon/grafana-dashboards-configmap-generator



monitoring/grafana/grafana-dashboards-configmap-generator/bin/grafana dashboards_generate.sh \







Grafana Watcher to reload Dashboards

1) Use 'grafana-watcher' container to reload Grafana dashboards supplied in ConfigMap volume

- name: grafana-watcher
image: <u>quay.io/coreos/grafana-watcher:v0.0.8</u>
args:
- 'watch-dir=/var/grafana-dashboards-0'
- 'grafana-url= <u>http://localhost:3000'</u>
env:
- name: GRAFANA_USER
valueFrom:
secretKeyRef:
name: grafana-credentials
key: user
 name: GRAFANA_PASSWORD
valueFrom:
secretKeyRef:
name: grafana-credentials
key: password
volumeMounts:
- name: grafana-dashboards-0
mountPath:/var/grafana-dashboards-0
volumes:
- name: grafana-storage
emptyDir: {}
- name: grafana-dashboards-0
configMap:
name: grafana-dashboards-0





- Control over which dashboards are deployed; some Grafana instances have dashboards that are unused or point to nonexistent servers

 Replace 'grafana-watcher' pod with Grafana provider config that automatically reloads dashboards from ConfigMap volume path





Credit to the Mux Team



Adam Brown



Matt Ward









