From Containers to Kubernetes Operators

Philipp Krenn @xeraa

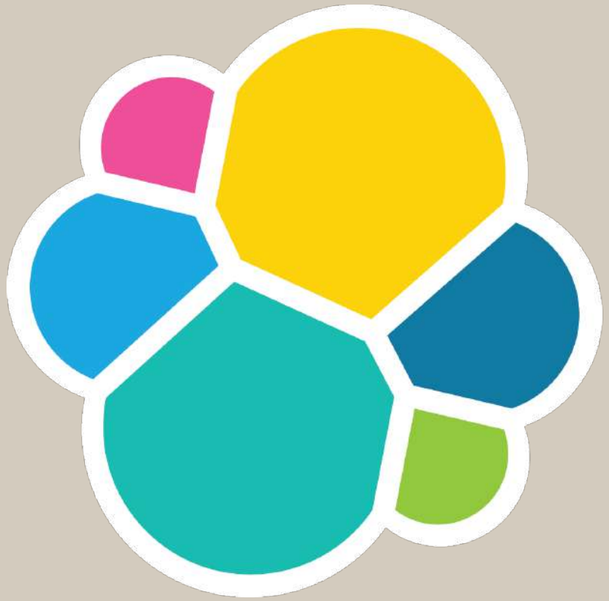

elastic

Developer 🥑

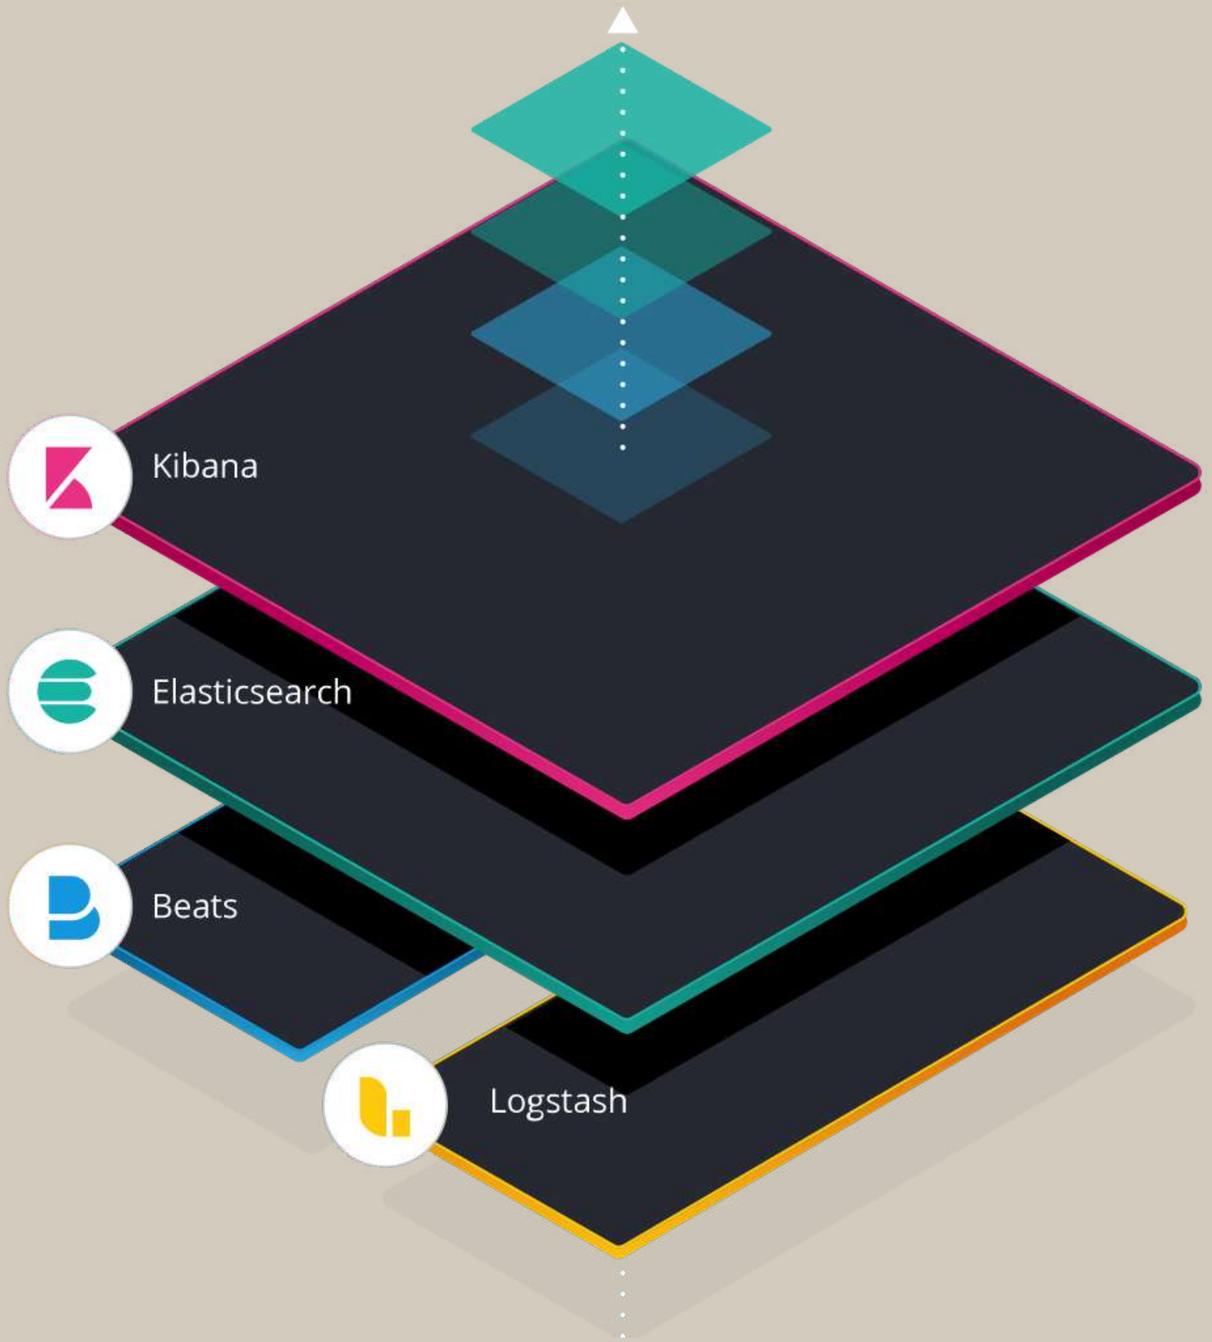
Kibana
Elasticsearch
Beats
Logstash

Who uses containers?

Who uses containers in production?

Who uses stateful containers?

Who uses containers from Elastic?

Agenda

Docker images
Docker Compose
Helm Chart
Kubernetes Operator

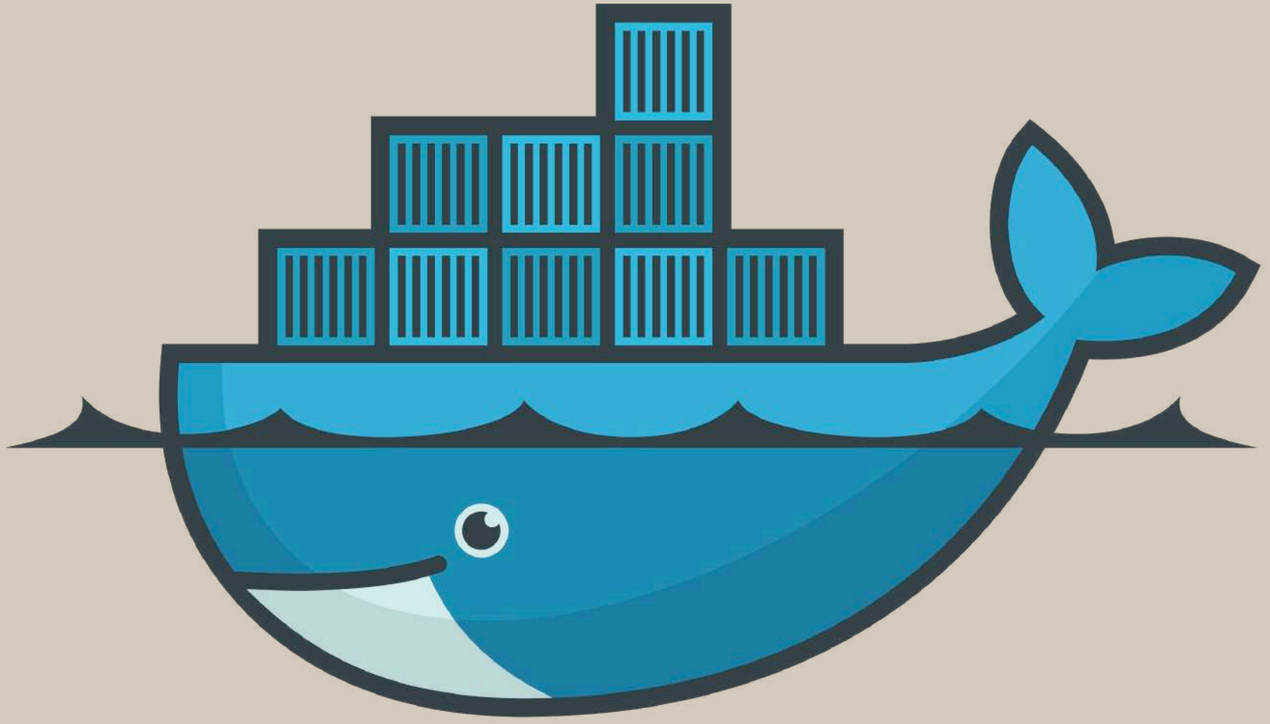

docker


HD

Docker: the world's most heavily funded college project

— Internal quote from Slack

Containers are the new ZIP format to distribute software

One of many…

RPM, DEB, TAR.GZ, MSI
Ansible, Chef, Puppet

…but not without issues

Fallacy
root and chmod 777

:-p sed s/🐍/🐳/g #Docker

The container runs Elasticsearch as user elasticsearch using uid:gid 1000:0.

https://www.elastic.co/guide/en/elasticsearch/reference/current/docker.html

bluepuma77 commented on 4 Mar

Bug description

Starting elasticsearch:5.2.2 results in `Failed to created node environment`. This has been covered in #21, the recommendation was to set user permissions on the host system.

```
docker run --rm -v /tmp/elastic:/usr/share/elasticsearch/data docker.elastic.co/elasticsear
# breaks
```

I prefer simple fire & forget docker containers which don't need any preparation on the host system. Would it be an option to check the Elasticsearch image from Docker, that runs fine, without setting any permissions before, and replicate that behaviour in new `elasticsearch-docker`?

```
docker run --rm -v /tmp/elastic:/usr/share/elasticsearch/data elasticsearch:5.2.2
# runs
```

```
docker run --rm -v /tmp/elastic:/usr/share/elasticsea
# breaks
```

I prefer simple fire & forget docker containers which don't n
Would it be an option to check the Elasticsearch image from
any permissions before, and replicate that behaviour in new

```
docker run --rm -v /tmp/elastic:/usr/share/elasticsea
# runs
```

Those who do not understand Unix are condemned to reinvent it, poorly.

— Henry Spencer

Fallacy
:latest

akes data easy to explore.

Copy and paste to pull this image

```
docker pull elasticsearch
```

View Available Tags

:7 and :7.13?

Base image & JVM are dependencies

Future

add image version?

Fallacy
Runtime mutation

Feature Request

```
if [ -f /custom/user_init.sh ]
then
. /custom/user_init.sh
fi
```

Or customization through environment variables

Dockerfile

```
ARG ELASTIC_VERSION
FROM docker.elastic.co/elasticsearch/elasticsearch:${ELASTIC_VERSION}

RUN bin/elasticsearch-plugin install analysis-phonetic --batch

ARG ELASTIC_VERSION
RUN bin/elasticsearch-plugin install \
    https://github.com/spinscale/elasticsearch-ingest-langdetect/releases/download/
    ${ELASTIC_VERSION}.1/ingest-langdetect-${ELASTIC_VERSION}.1.zip --batch
```

Generate Keystore

```
$ docker run -p 9200:9200 -p 9300:9300 -e "discovery.type=single-node" \
  -v /Users/philipp/Desktop/demo/config/:/usr/share/elasticsearch/config/ \
  -it docker.elastic.co/elasticsearch/elasticsearch:7.13.2 /bin/bash
[root@1006ed50b646 elasticsearch]# ./bin/elasticsearch-keystore create
Created elasticsearch keystore in /usr/share/elasticsearch/config
[root@1006ed50b646 elasticsearch]# ./bin/elasticsearch-keystore add test
Enter value for test:
[root@1006ed50b646 elasticsearch]# exit
exit
$ cat config/elasticsearch.keystore
??lelasticsearch.keystore?@g?o!?$?K?Lf?w?VAEŠЉm?[?a6?B??? y?,!B}??Ħ?æ?AU=?C?:?o?
?W?O8?}U?;p?Ţ???cQ?????7?JY?     2A?:???ZUY??2V?9?ᘐ??(??0?q\
```

Mount Keystore

Docker Compose

```
elasticsearch:
    image: docker.elastic.co/elasticsearch/elasticsearch:${ELASTIC_VERSION}
    secrets:
      - source: elasticsearch.keystore
        target: /usr/share/elasticsearch/config/elasticsearch.keystore
```

Fallacy
Base image

Common base image since 5.4+

CentOS 7

Upside

Similar setup

Shared layers

Upside

JVM large anyway

glibc tested

Downside

Size

Does it matter?

Stateful vs Stateless

ARM64

Supported

Slimmed Down CentOS 8

Version 8.0

https://github.com/elastic/elasticsearch/pull/52519

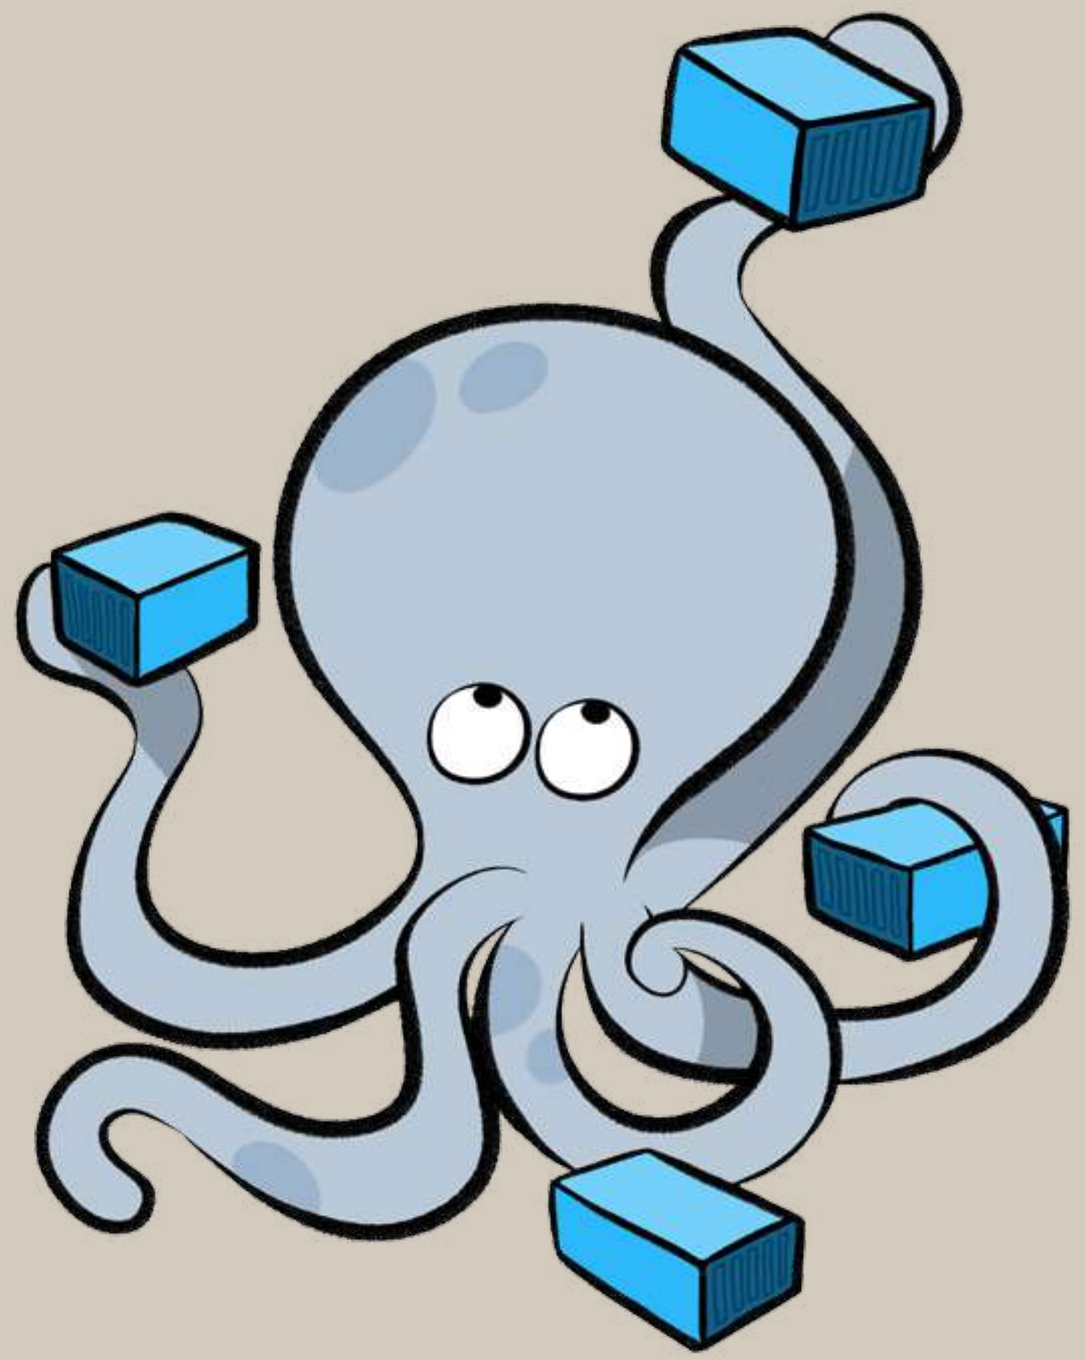

Who uses

Docker Compose?

More than

development / demo?

```yaml
---
version: '2'
services:

  elasticsearch:
    image: docker.elastic.co/elasticsearch/elasticsearch:$ELASTIC_VERSION
    environment:
      - "ES_JAVA_OPTS=-Xms512m -Xmx512m"
      - discovery.type=single-node
    mem_limit: 1g
    volumes:
      - esdata:/usr/share/elasticsearch/data
    ports:
      - 9200:9200

  kibana:
    image: docker.elastic.co/kibana/kibana:$ELASTIC_VERSION
    links:
      - elasticsearch
    ports:
      - 5601:5601

volumes:
  esdata:
    driver: local
```

Deprecated

https://github.com/elastic/stack-docker

kubernetes

Who uses

Kubernetes?

Kubernetes is the answer. What was the question?

— **https://twitter.com/charlesfitz/status/1068203930683752448**

Who uses

Swarm?

Who uses

Nomad?

Who uses

Mesos?

Static configuration to deploy resources to K8s

Kubernetes Architecture

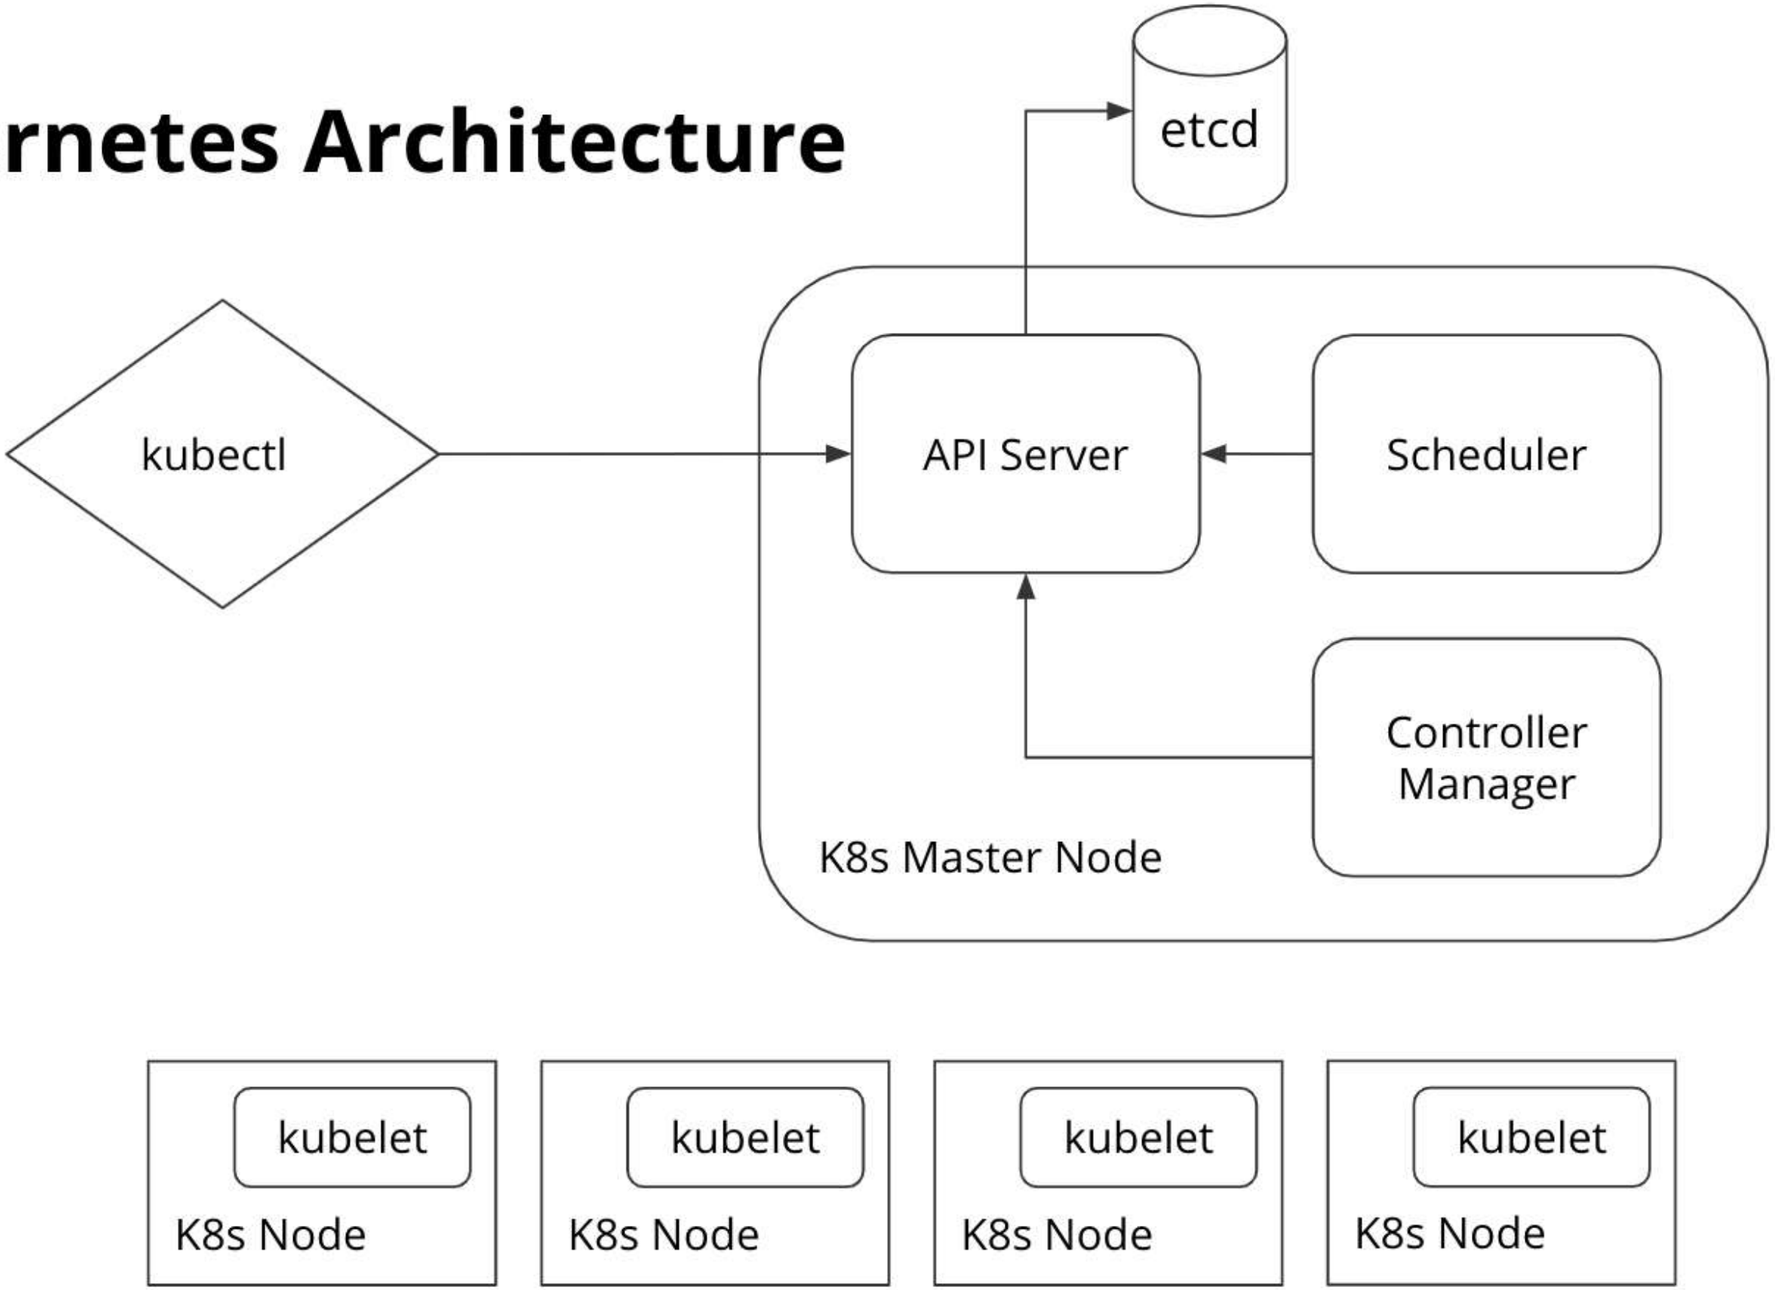


YA
ML

…lots of it

Fun with YAML

http://www.yamllint.com

```yaml
ports:
  - 80:80
  - 20:20
```

Fun with YAML

https://docs.docker.com/compose/compose-file/#short-syntax-1

```yaml
ports:
  - "80:80"
  - 73200
```

Kubernetes 1.8 allows dots in env vars

https://github.com/kubernetes/kubernetes/issues/2707


HELM

Advanced package management with support for templating and more complex resources

Building on existing Kubernetes primitives like StatefulSet, Service, Deployment,…

Elastic Helm Charts

Elasticsearch, Kibana, Filebeat, Metricbeat, APM Server, Logstash

https://github.com/elastic/helm-charts

StatefulSet (STS)

By default rolling upgrade

Waiting for cluster health being green after each instance is updated

Un-Opinionated

Expose environment variables & mount secrets

Multiple upgrade strategies

Example

https://github.com/elastic/helm-charts/tree/master/elasticsearch/examples

```
helm repo add elastic https://helm.elastic.co

helm repo update

helm install elasticsearch elastic/elasticsearch [--set imageTag=7.13.2]

cd examples/default
make
```

```yaml
---
# Shrink default JVM heap
esJavaOpts: "-Xmx128m -Xms128m"

# Allocate smaller chunks of memory per pod
resources:
  requests:
    cpu: "100m"
    memory: "512M"
  limits:
    cpu: "1000m"
    memory: "512M"

# Request smaller persistent volumes
volumeClaimTemplate:
  accessModes: [ "ReadWriteOnce" ]
  storageClassName: "standard"
  resources:
    requests:
      storage: 100M
```

OPERATOR
LIFECYCLE MANAGER

Expands K8s interfaces and allows to customize management of applications

Custom Resource Definition (CRD)

Think: Elasticsearch, Kibana, APM

Contrast: Built-in resources like Pods, Services, Secrets, StatefulSets,…

Custom Resource (CR)

CRD == type definition (class)

CR == instance (object)

Custom Controller

Brings CRDs to "life" (reconciliation loop)

Upgrades, secrets, certificate management,…

Elastic Operator

Elasticsearch, Kibana, Beats / Agent, APM Server, Enterprise Search

https://github.com/elastic/cloud-on-k8s

Golang 1.13+

Kubebuilder 2

SDK for building Kubernetes APIs using CRDs

Kustomize

Generate patched CRDs for specific flavors

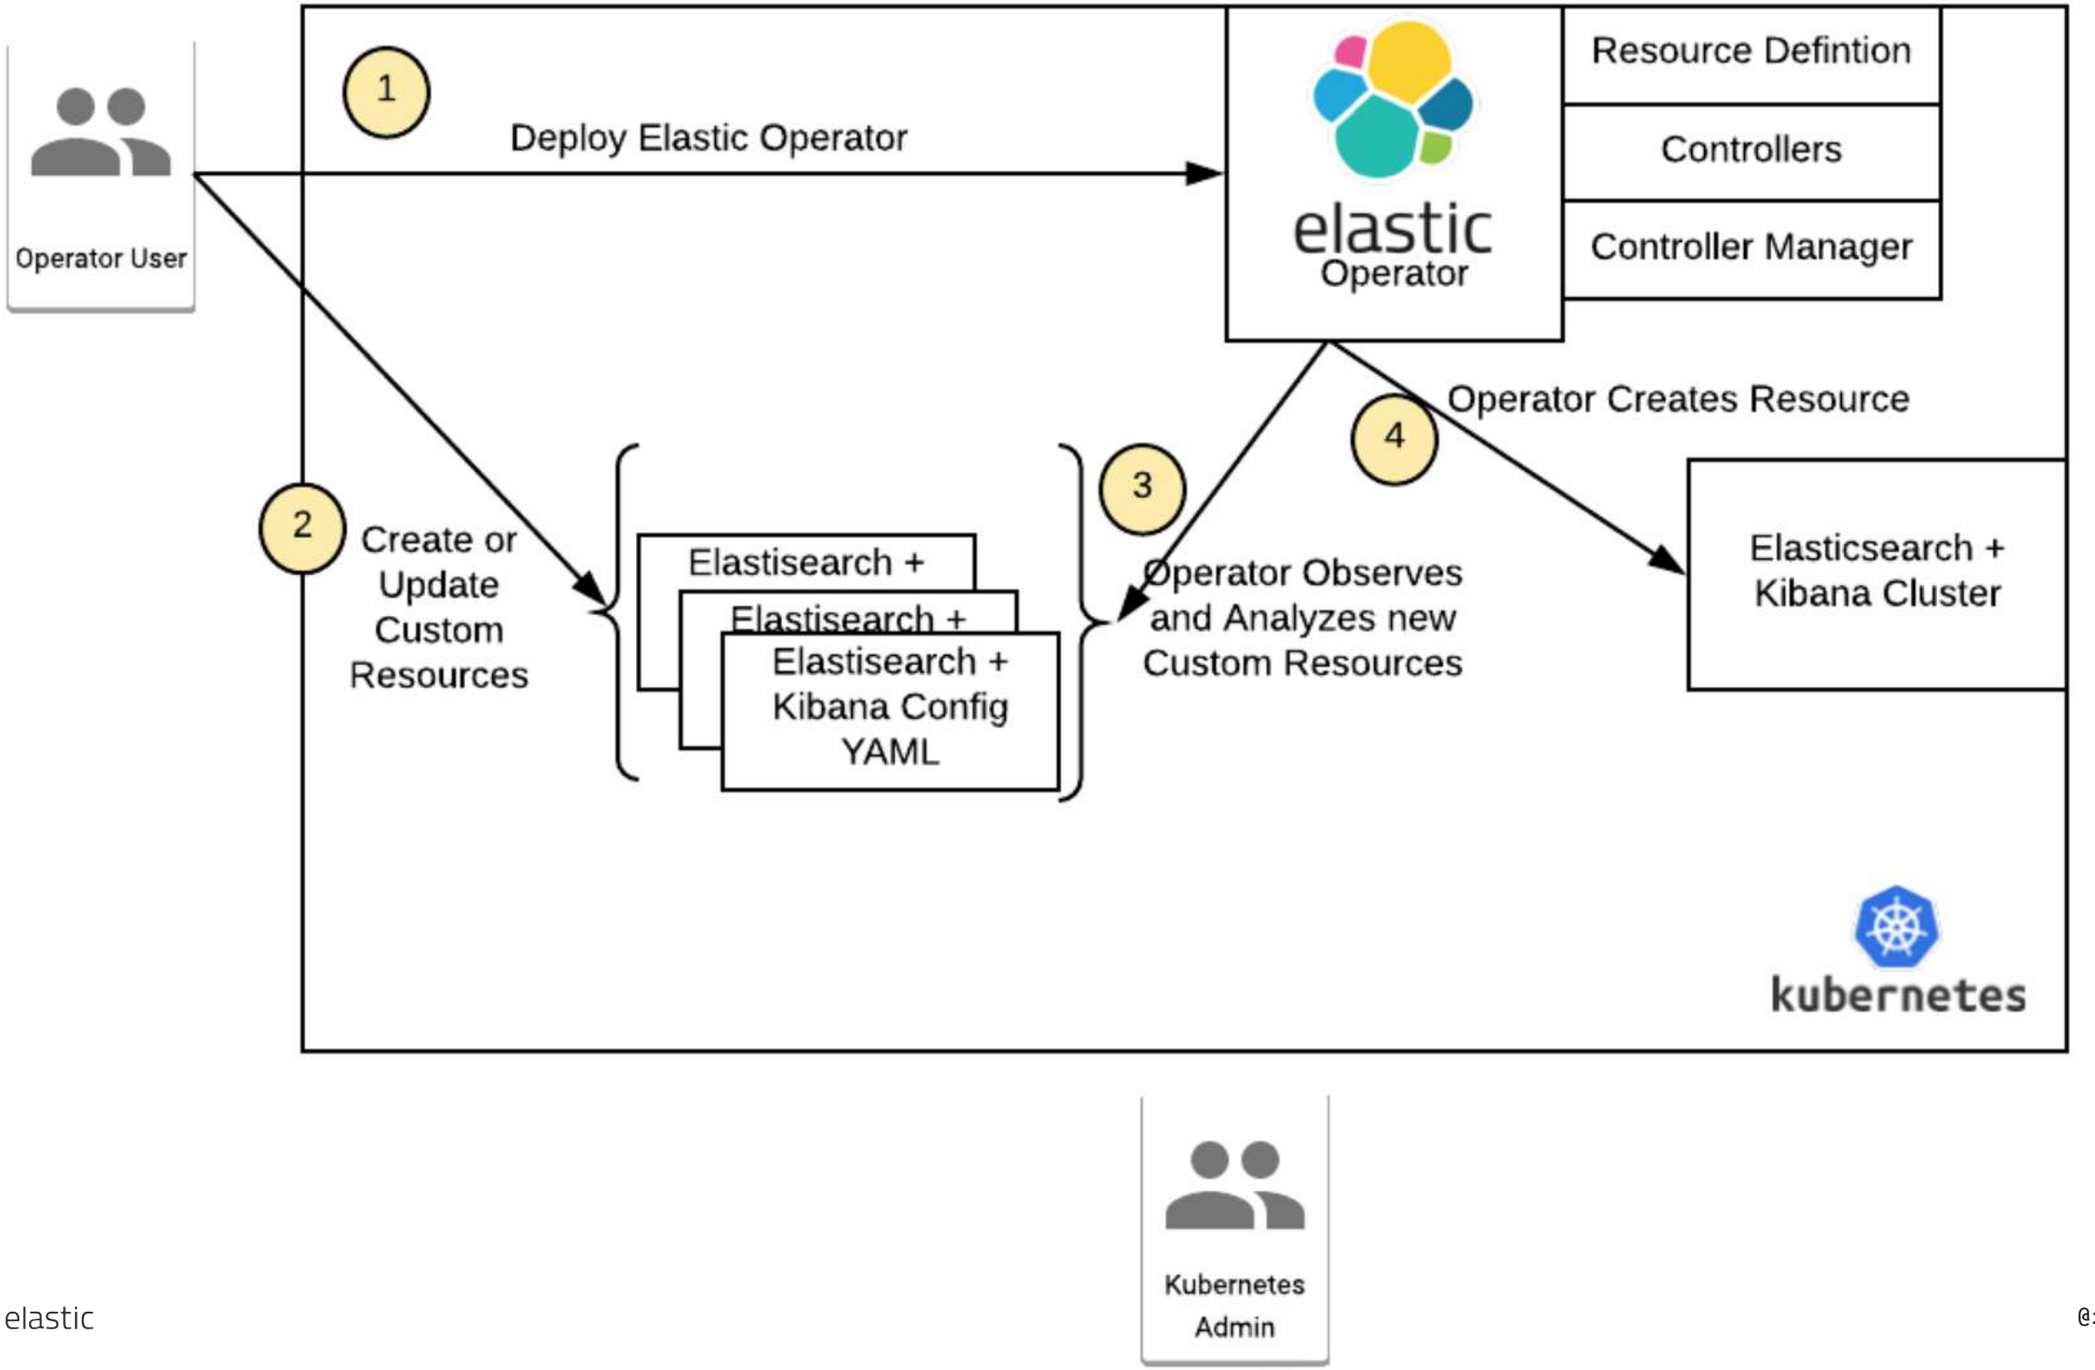
Operator User
1
Deploy Elastic Operator
elastic
Operator
Resource Defintion
Controllers
Controller Manager
Operator Creates Resource
4
2
Create or
Update
Custom
Resources
Elastisearch +
Elastisearch +
Elastisearch +
Kibana Config
YAML
3
Operator Observes
and Analyzes new
Custom Resources
Elasticsearch +
Kibana Cluster
kubernetes
Kubernetes
Admin

Opinionated

Encode best practices & operational knowledge
Built-in certificate management, security,…

Example Opinions

Scale down: Drain nodes first

Upgrade: Disable shard allocation

You Can Still Shoot Yourself in the Foot

Configure 0 replicas and do an upgrade for example

Quick Setup

```
# Get the current version
kubectl apply -f https://download.elastic.co/downloads/eck/1.6.0/all-in-one.yaml

# Monitor logs
kubectl -n elastic-system logs -f statefulset.apps/elastic-operator

# Deploy
kubectl apply -f es_kibana.yml
```

Helm Chart Setup

```
helm repo add elastic https://helm.elastic.co

helm repo update

helm install elastic-operator elastic/eck-operator \
     -n elastic-system --create-namespace

kubectl -n elastic-system logs -f statefulset.apps/elastic-operator

kubectl get pods --all-namespaces
```

```yaml
---
apiVersion: elasticsearch.k8s.elastic.co/v1
kind: Elasticsearch
metadata:
  name: elasticsearch-sample
spec:
  version: 7.13.2
  nodeSets:
  - name: default
    count: 1
    config:
      node.store.allow_mmap: false
```

```yaml
---
apiVersion: kibana.k8s.elastic.co/v1
kind: Kibana
metadata:
  name: kibana-sample
spec:
  version: 7.13.2
  count: 1
  elasticsearchRef:
    name: "elasticsearch-sample"
  http:
    service:
      spec:
        type: LoadBalancer
```

Access Kibana

```
# Check status
kubectl get elasticsearch,kibana
# Get the credentials
echo `kubectl get secret elasticsearch-sample-es-elastic-user
     -o=jsonpath='{.data.elastic}' | base64 --decode`
```

Changes

Instance size / number, version,…

kubectl apply -f es_kibana.yml

System Settings

node.store.allow_mmap: false
(for demos only)

Init Container
(privileged security context)

Manual setting on the host
(recommended)

```yaml
apiVersion: elasticsearch.k8s.elastic.co/v1
kind: Elasticsearch
metadata:
  name: elasticsearch-sample
spec:
  version: 7.13.2
  nodeSets:
  - name: default
    count: 3
    podTemplate:
      spec:
        initContainers:
        - name: sysctl
          securityContext:
            privileged: true
          command: ['sh', '-c', 'sysctl -w vm.max_map_count=262144']
```

Support

GKE (Google Cloud)

EKS (AWS)

AKS (Azure)

OpenShift (Redhat)

StatefulSets (STS)

Rolling Upgrades with Volume reuse

"Standard" way to run stateful workloads — stable network ID, stable data volume that is re-attachable during rolling upgrades

Storage (1)

Persistent Volumes (default)

Local PV, Cloud vendor specific network attached

Storage (2)

EmptyDir / HostPath

Not durable, no enforceable quota, may lead to eviction

Global Namespace

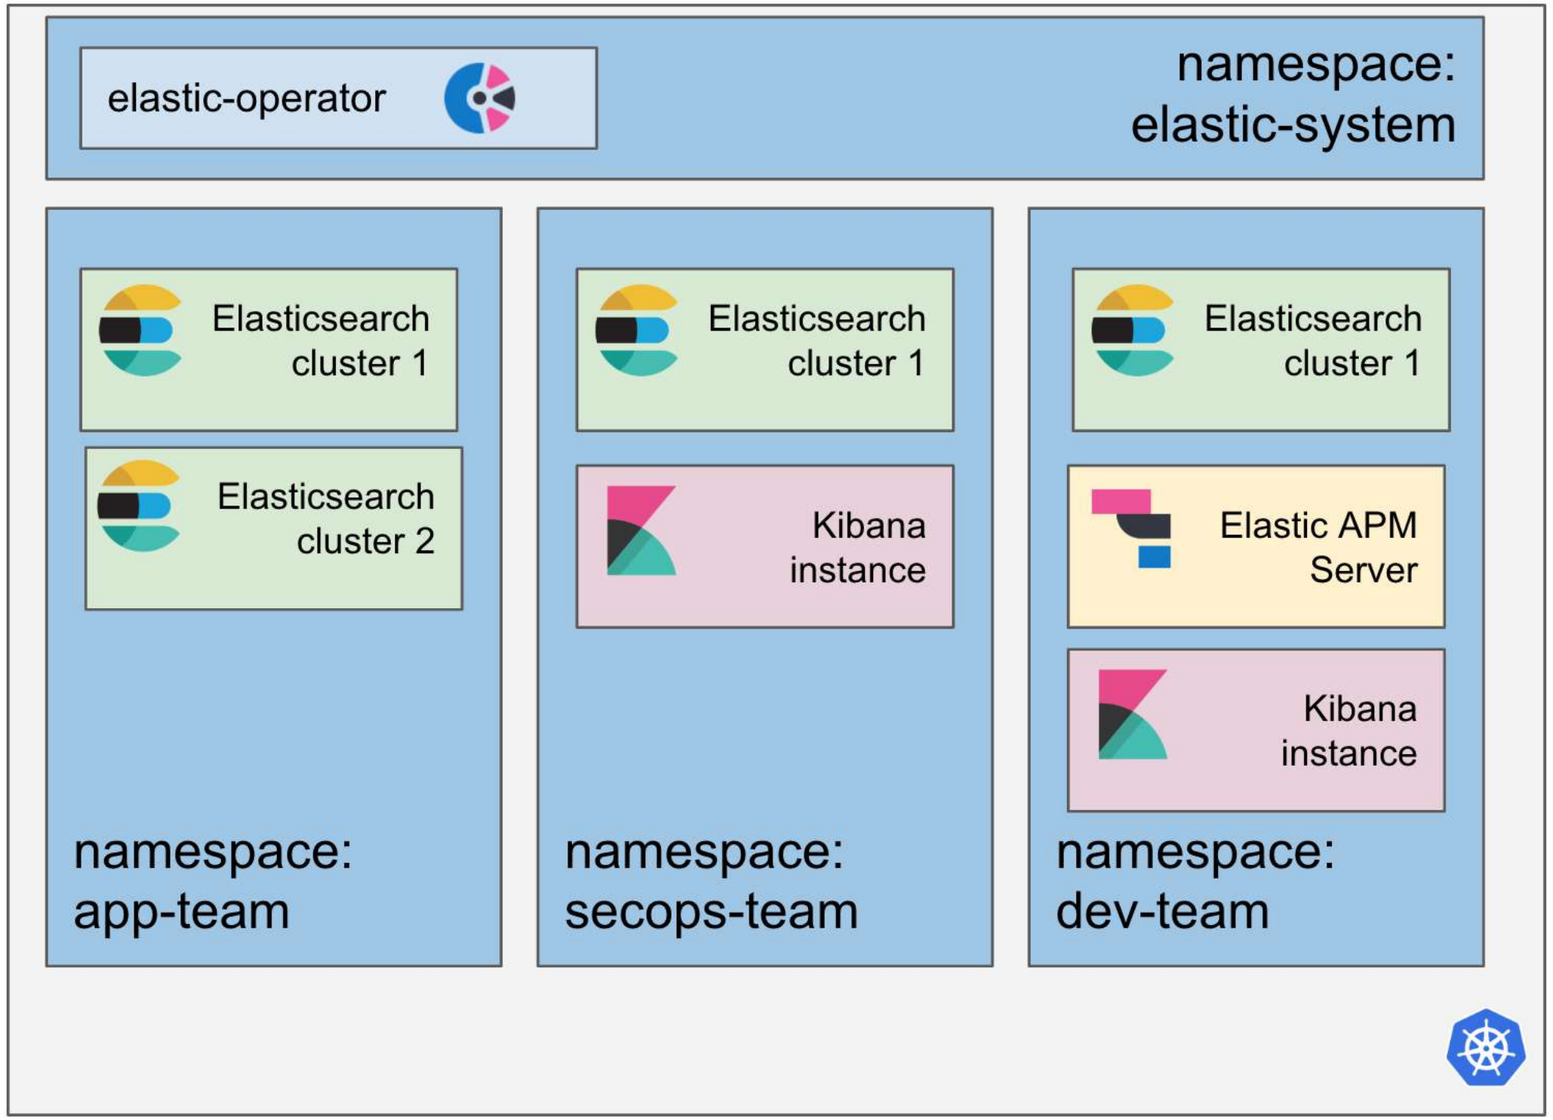

Single Namespace

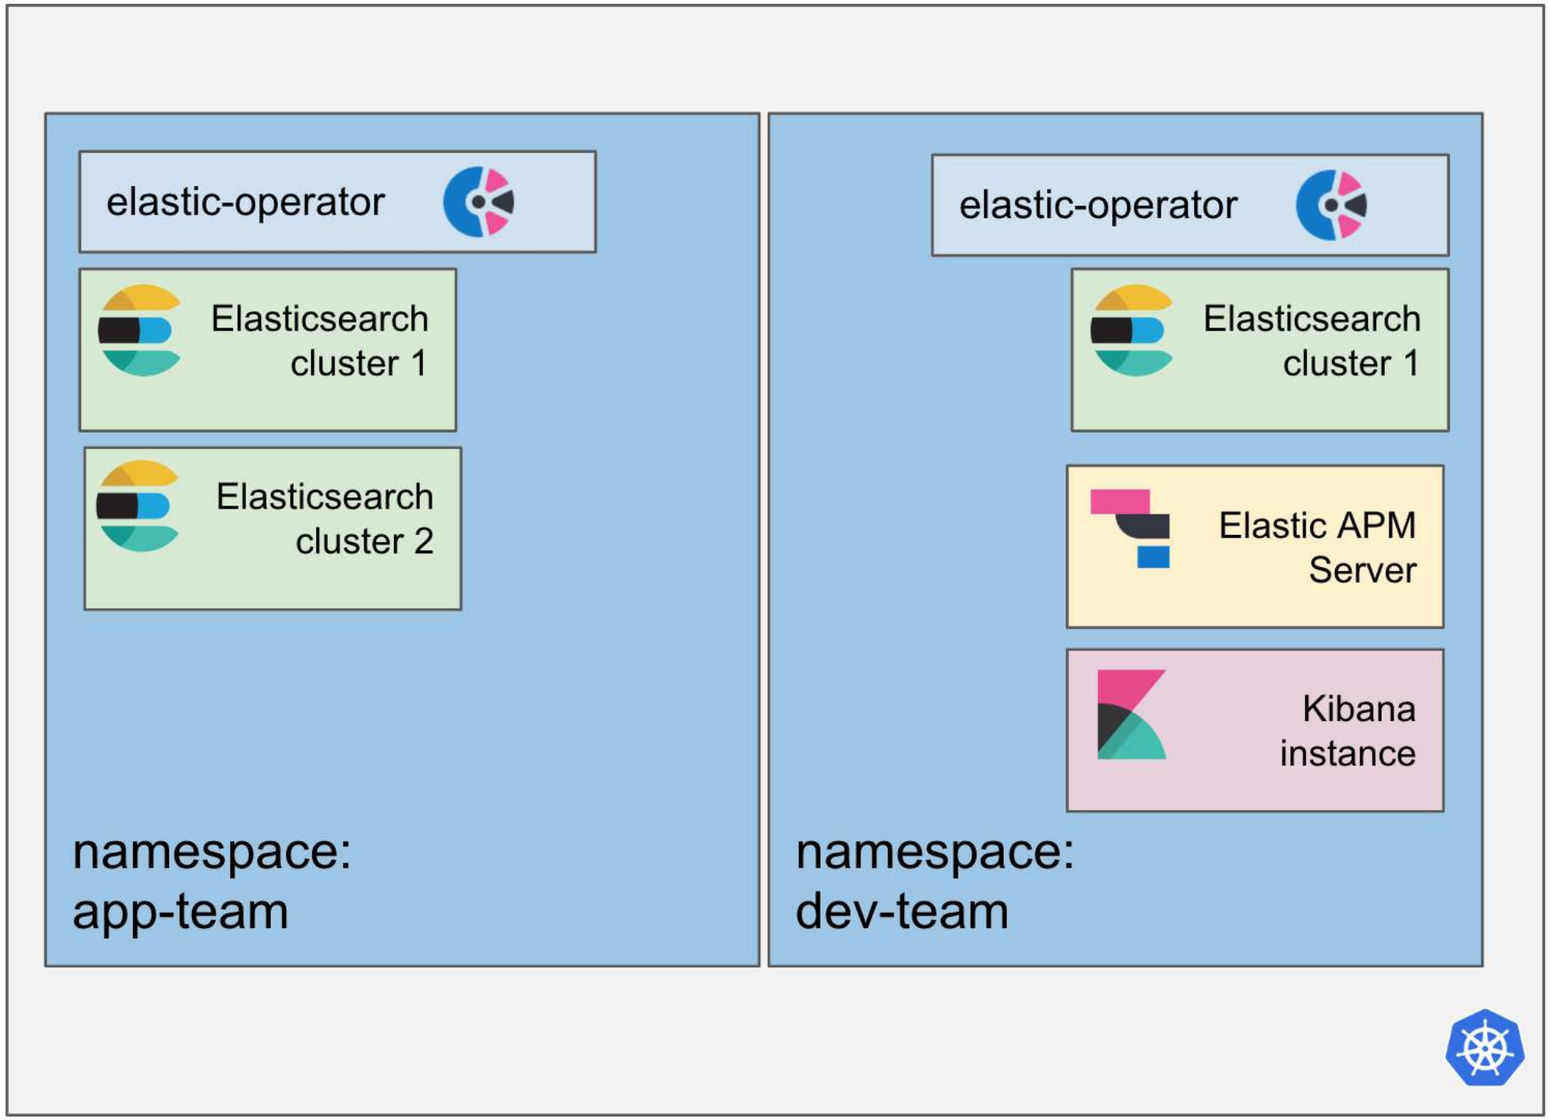

Other Operators

MongoDB, Kafka, Redis, CockroachDB,…

Operator "Marketplaces": https://operatorhub.io

Couchbase Operator

provided by Couchbase

The Couchbase Autonomous Operator allows users to easily deploy, manage, and maint

Crunchy PostgreSQL Enterprise

provided by CrunchyData.com

A Postgres Operator from Crunchydata.com

Dynatrace OneAgent

provided by Dynatrace LLC

Install full-stack monitoring of Kubernetes clusters with the Dynatrace OneAgent.

Eclipse Che

provided by Eclipse Foundation

A Kube-native development solution that delivers porta

Elastic Cloud on Kubernetes

provided by Elastic

Run Elasticsearch, Kibana and the APM Server on Kubernetes and OpenShift

EnMasse

provided by EnMasse

EnMasse provides messaging as a managed service on Kubernetes

etcd

provided by CNCF

Create and maintain highly-available etcd clusters on Kubernetes

Event Streams Topic

provided by IBM

An operator for managing Topics for Event Streams on IBM Cloud

Ext Postgres Operator

provided by movetokube.com

Manage databases and roles in external PostgreSQL server or

Falco Operator

provided by Sysdig

Falco is a behavioral activity monitor designed to detect

Federator.ai

provided by ProphetStor Data Services, Inc.

Federator.ai Operator provides

FfDL Operator

provided by IBM

Fabric for Deep Learning - an operating system fabric for

Conclusion

"Containers are disrupting the industry!"

"Can I run Elasticsearch on Docker or Kubernetes?"

"Should I run Elasticsearch on Docker or Kubernetes?"

Kubernetes Paradox

"You don't have an Operator? That's a blocker!"

Kubernetes Paradox

"So you run the majority of your apps on Kubernetes?"

"Actually, we are just starting with Kubernetes and have 2% of our workloads on it."

Helm Charts vs Operator

Day 2 Operations

Questions & Discussion

Philipp Krenn @xeraa