


# From Containers to Kubernetes Operators

Emanuil Tolev

@emanuil\_tolev

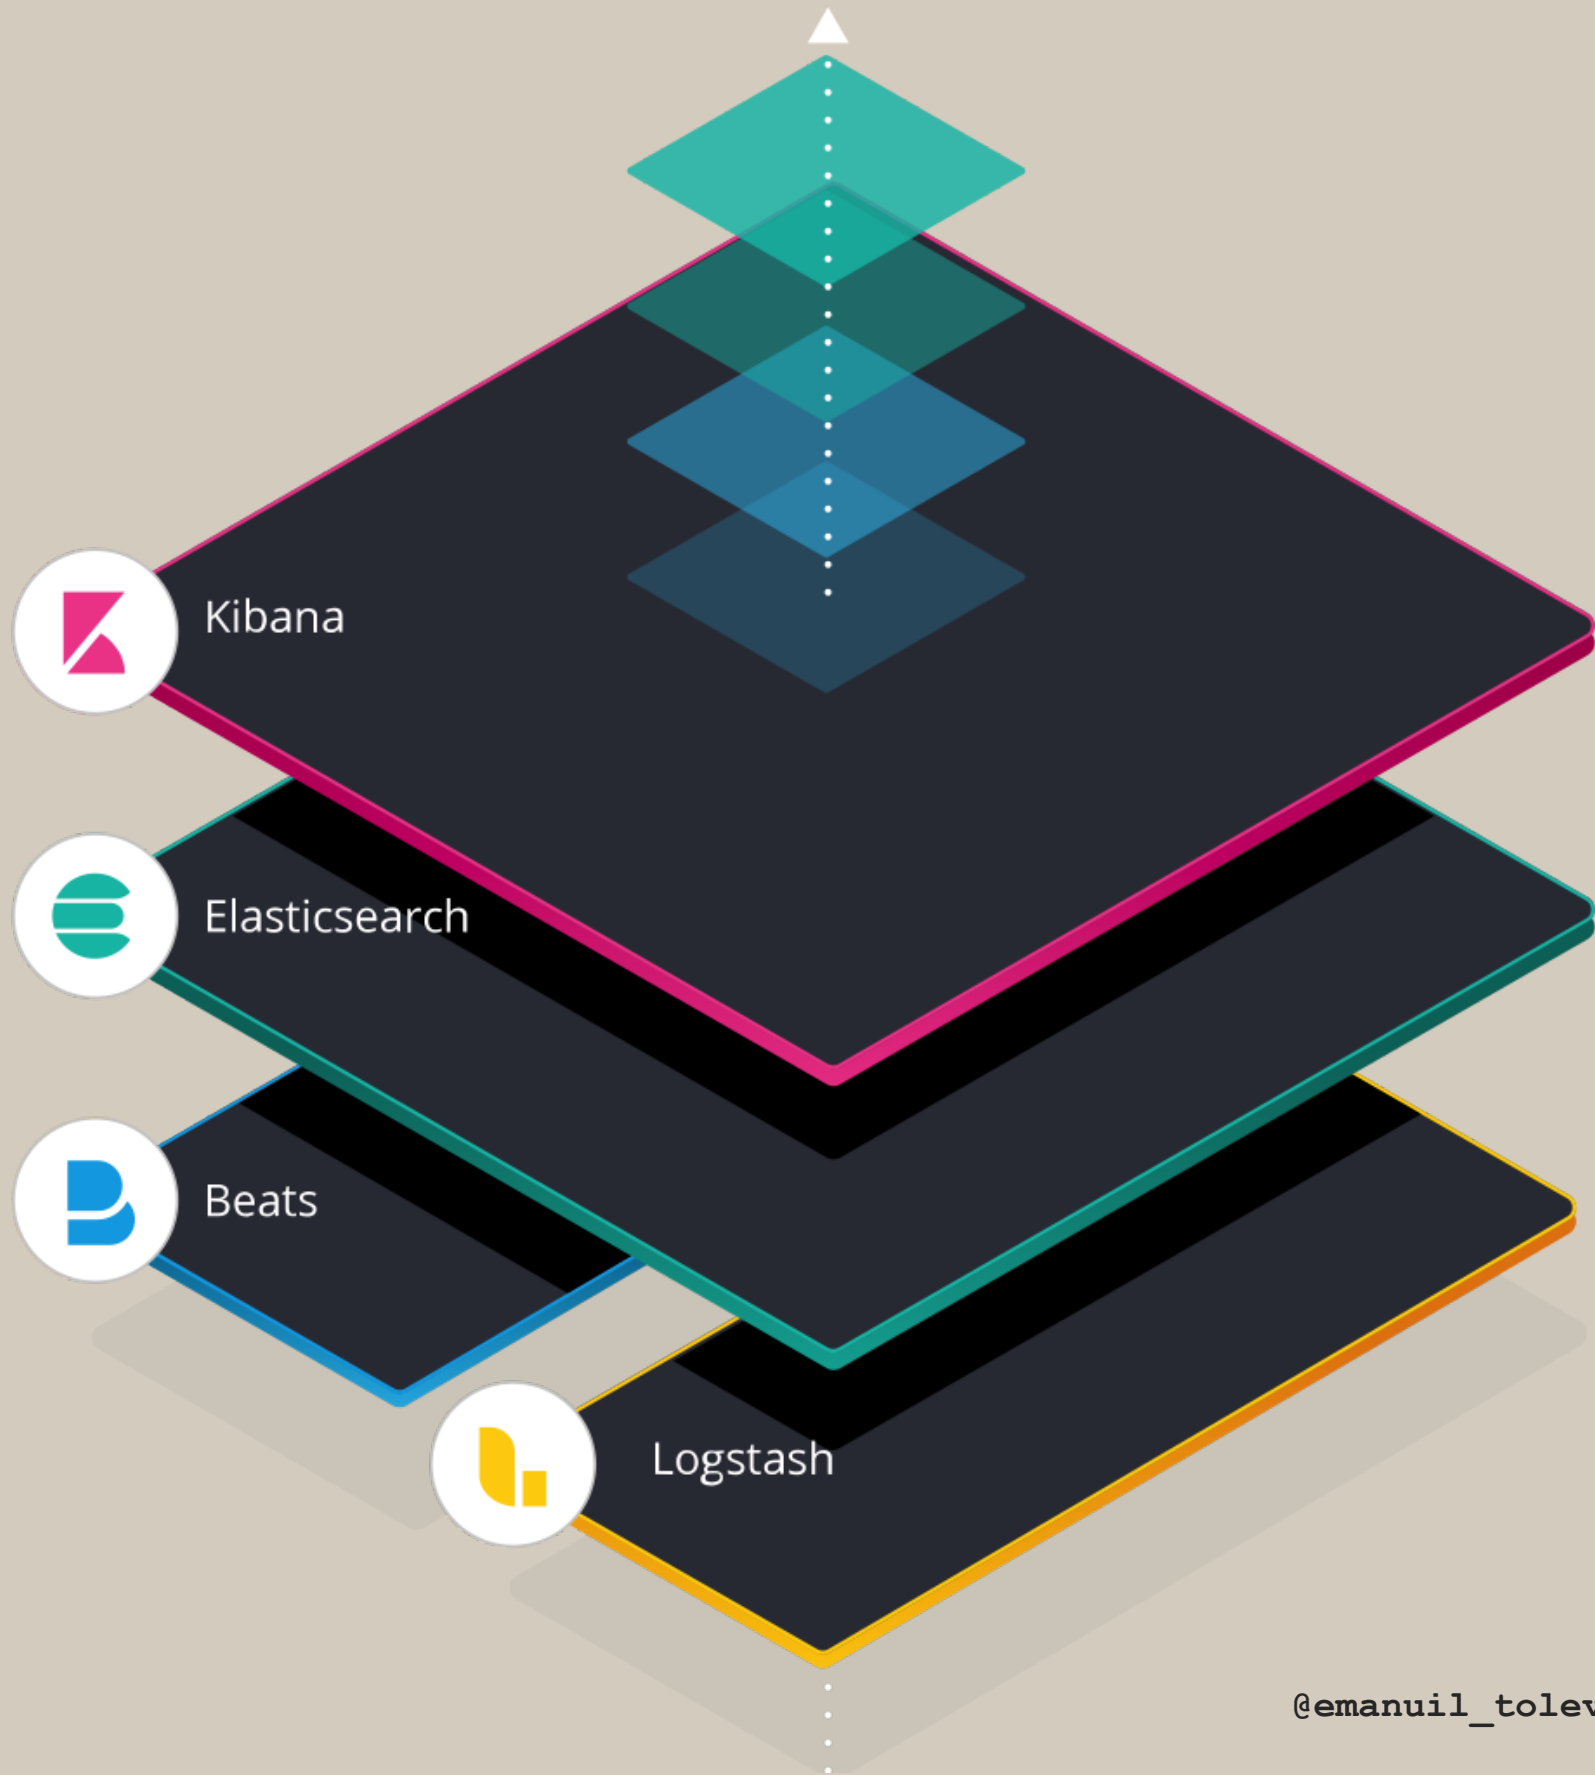


@emanuil\_tolev



elastic

Community Engineer



# Agenda - Elastic's journey through

Docker images

Helm Chart

Kubernetes Operator







**Containers are the new ZIP  
format to distribute software**

# One of many...

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

Fallacy  
**: latest**



makes data easy to explore.

Copy and paste to pull this image

```
docker pull elasticsearch
```



[View Available Tags](#)

No `:latest`, what  
about `:7` and `:7.7`?

Exact only: `7.7.1`

# The base image diversity and size debate

Common base image since 5.4+

CentOS 7

Similar setup

Shared layers - ultimately  
much better for size across  
components





# kubernetes



@emanuil\_tolev

*Kubernetes is the answer. What was the question?*

– <https://twitter.com/charlesfitz/status/1068203930683752448>

**Y A**  
**M L**

...lots of it

# Fun with YAML

<http://www.yamllint.com>

```
ports:
```

- 80:80
- 20:20

# Fun with YAML

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

ports:

- "80:80"

- 73200 # 20\*60(1):20\*60(0)=1220 \* 60=73200



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>

# Tested on GKE

Default storage pd-ssd (network attached)

Kubernetes  $\geq 1.10$  supports Local  
PersistentVolumes for increased performance

# Un-Opinionated

Expose environment variables & mount secrets

Multiple upgrade strategies

# Minikube Example

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

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

```
helm install --name elasticsearch elastic/elasticsearch [--set imageTag=7.7.1]
```

```
minikube addons enable default-storageclass
```

```
minikube addons enable storage-provisioner
```

```
cd examples/minikube
```

```
make
```



---

# Permit co-located instances for solitary minikube virtual machines

antiAffinity: "soft"

# 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 FRAMEWORK

# Custom Resource (CR)

CRD == type definition (class)

CR == instance (object)

# Custom Resource Definition (CRD)

Think: Elasticsearch, Kibana, APM

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



# Custom Controller

Brings CRDs to "life" (reconciliation loop)  
Upgrades, secrets, certificate management,...



# Elastic Operator

Elasticsearch, Kibana,  
APM Server

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

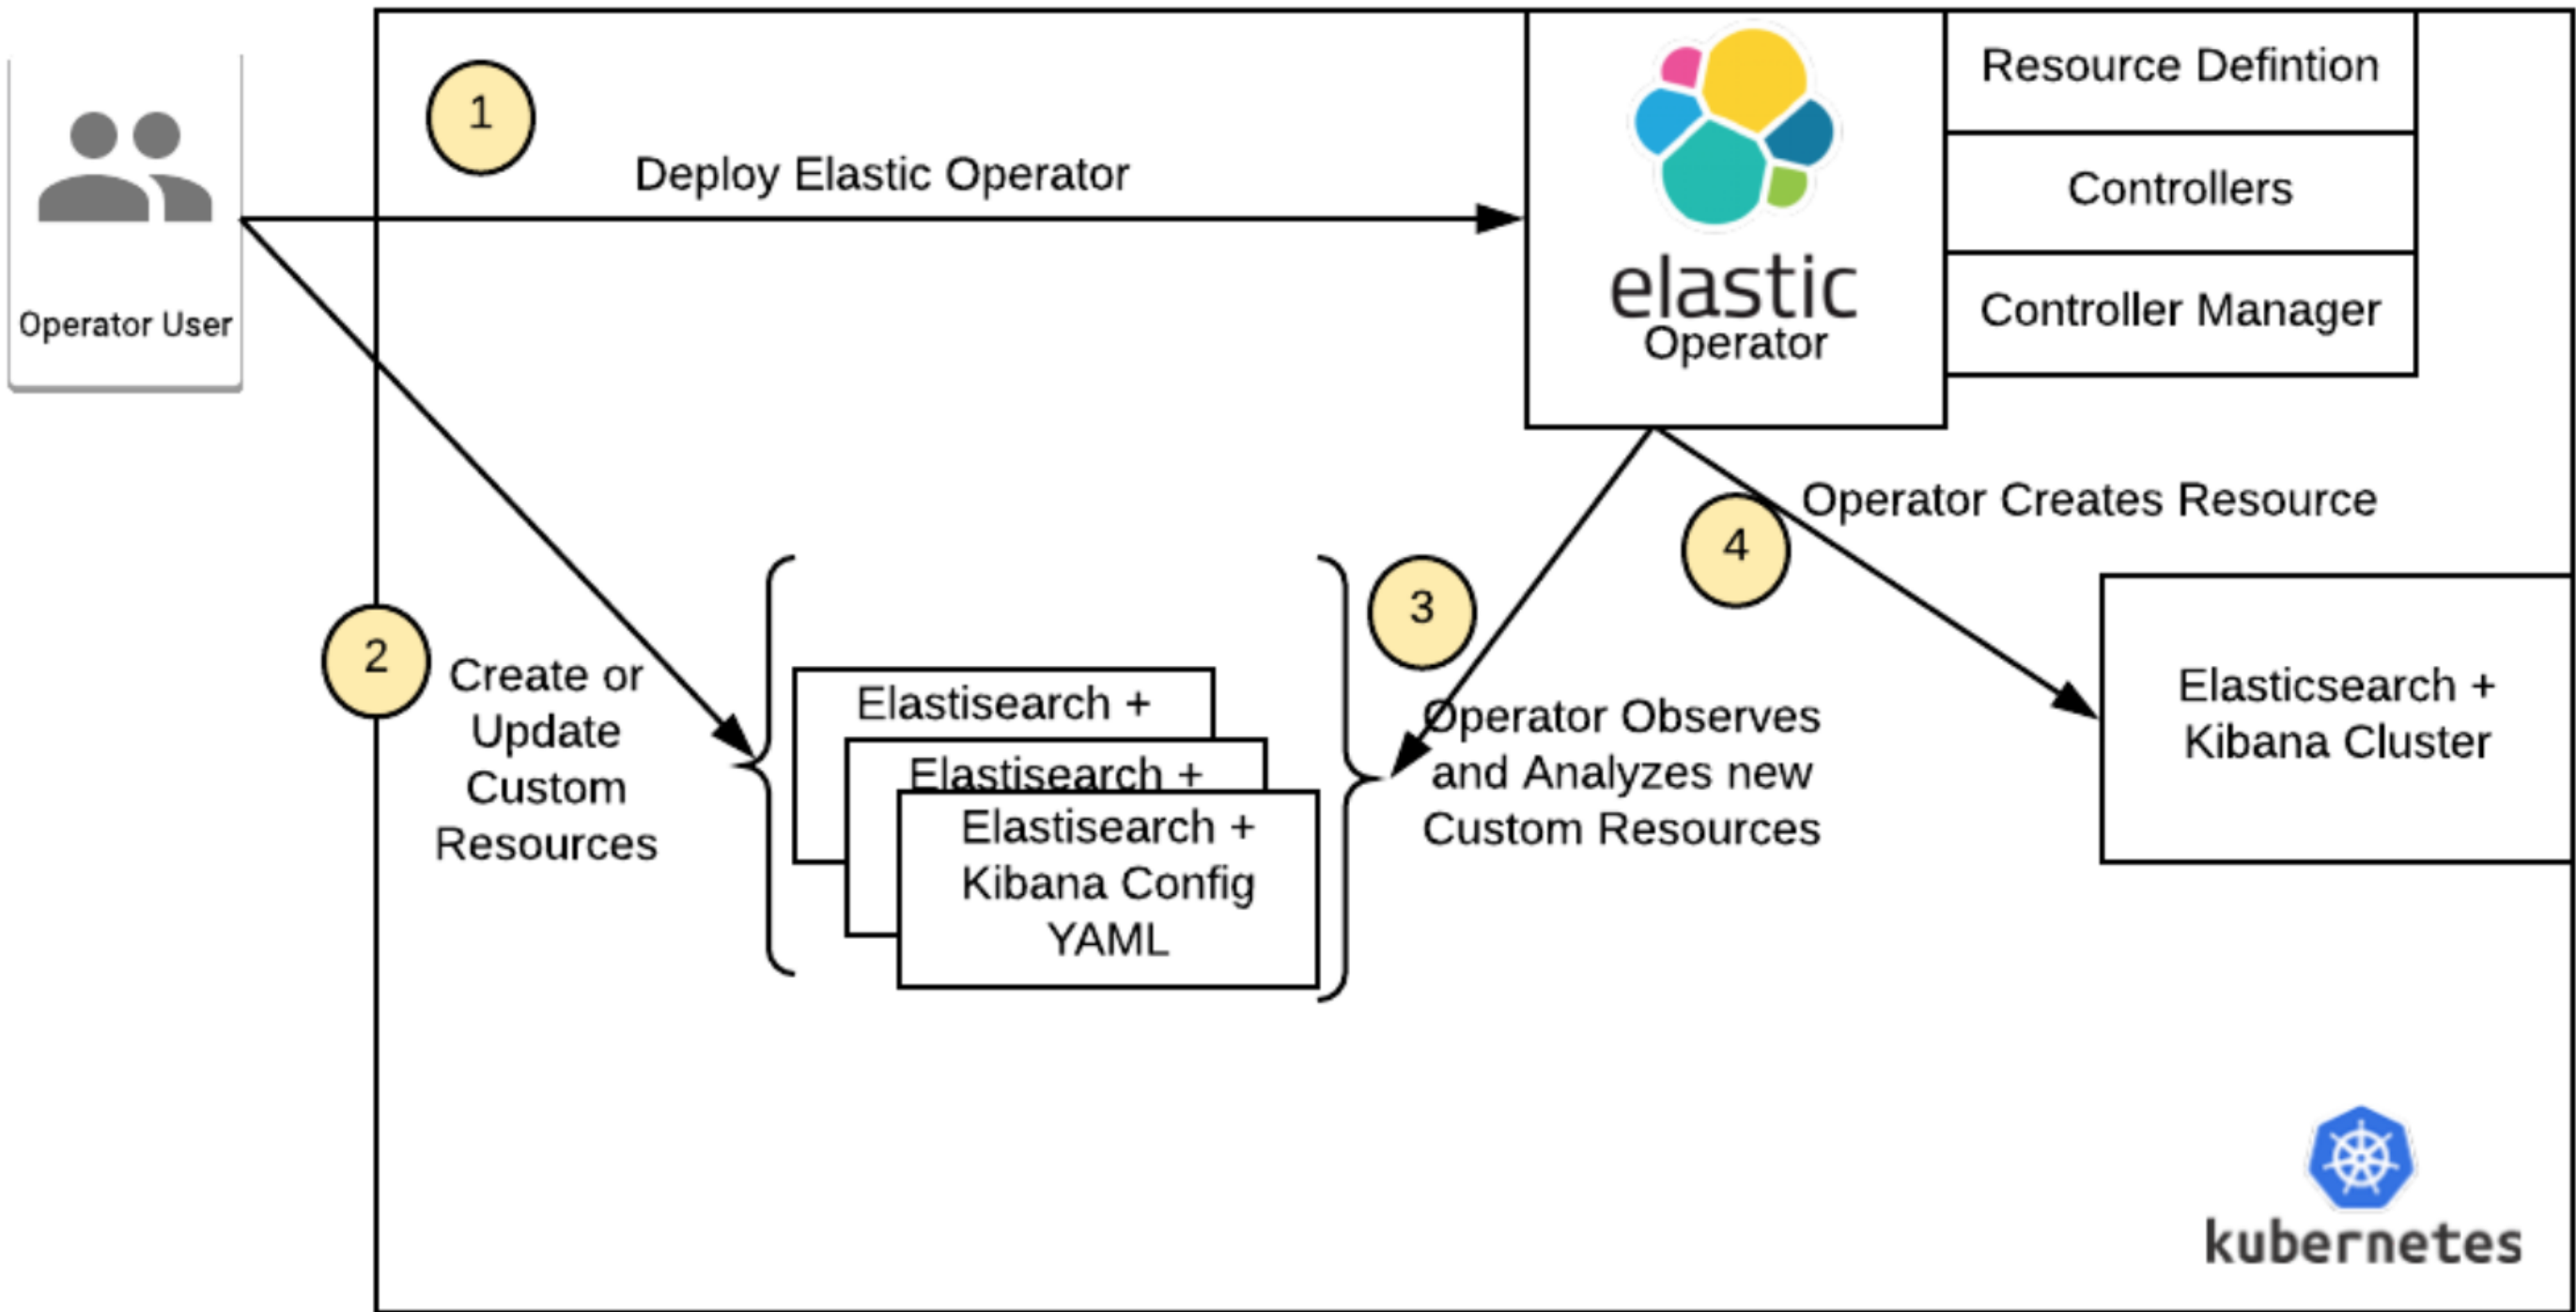
# Golang 1.13

# Kubebuilder 2

SDK for building Kubernetes APIs using CRDs

# Kustomize

Generate patched CRDs for specific flavors



# 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



# Running on Minikube

```
minikube config set memory 16384
```

```
minikube config set cpus 4
```

```
minikube start
```

# Running on Minikube

```
# Set up the entire operator: configs, deployment practices, monitoring, in one command  
kubectl apply -f https://download.elastic.co/downloads/eck/1.1.2/all-in-one.yaml
```

```
# Monitor logs
```

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

```
# And this is where you come in - the configs you write
```

```
kubectl apply -f apm_es_kibana.yaml
```



```
---
apiVersion: elasticsearch.k8s.elastic.co/v1
kind: Elasticsearch
metadata:
  name: elasticsearch-sample
spec:
  version: 7.7.1
  nodes:
  - nodeCount: 1
    podTemplate:
      spec:
        containers:
        - name: elasticsearch
          resources:
            limits:
              memory: 2Gi
  volumeClaimTemplates:
  - metadata:
      name: data
    spec:
      accessModes:
      - ReadWriteOnce
      resources:
        requests:
          storage: 2Gi
```

---

apiVersion: apm.k8s.elastic.co/v1

kind: ApmServer

metadata:

name: apm-server-sample

spec:

version: 7.7.1

nodeCount: 1

elasticsearchRef:

name: "elasticsearch-sample"



---

```
apiVersion: kibana.k8s.elastic.co/v1
```

```
kind: Kibana
```

```
metadata:
```

```
  name: kibana-sample
```

```
spec:
```

```
  version: 7.7.1
```

```
  nodeCount: 1
```

```
  elasticsearchRef:
```

```
    name: "elasticsearch-sample"
```



# Running on Minikube

```
# Check status
```

```
kubectl get elasticsearch,kibana,apmserver
```

```
# Expose Kibana
```

```
kubectl port-forward service/kibana-sample-kb-http 5601
```

```
# 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 apm_es_kibana.yaml
```

# Support

GKE (Google Cloud)

EKS (AWS)

AKS (Azure)

OpenShift (Redhat)



# StatefulSets

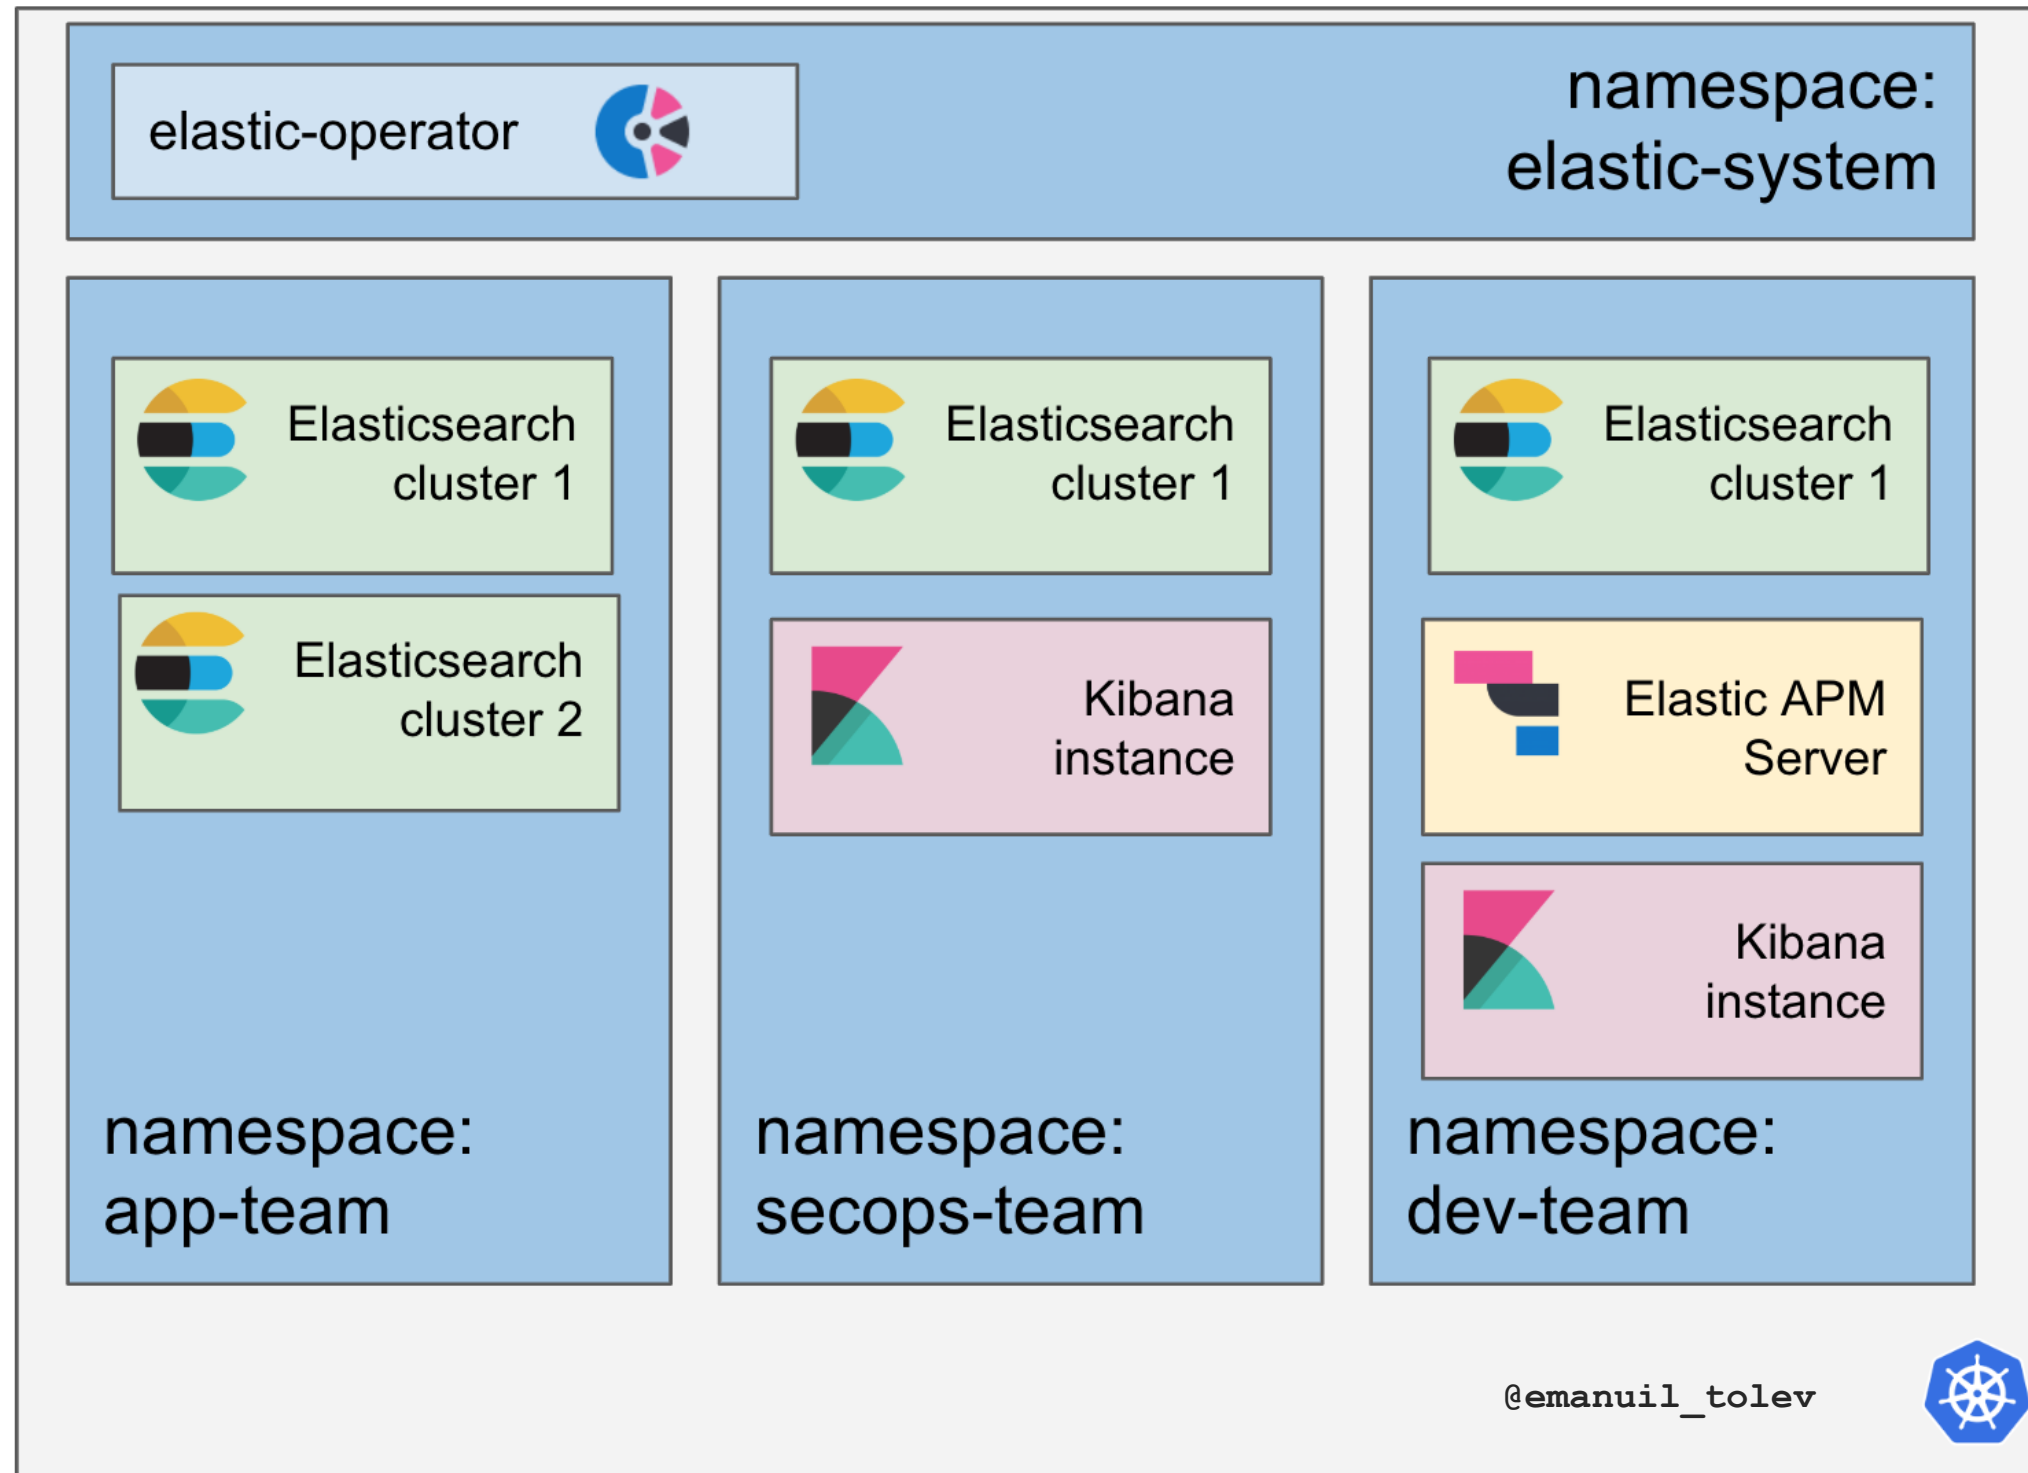
Rolling Upgrades with Volume reuse

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

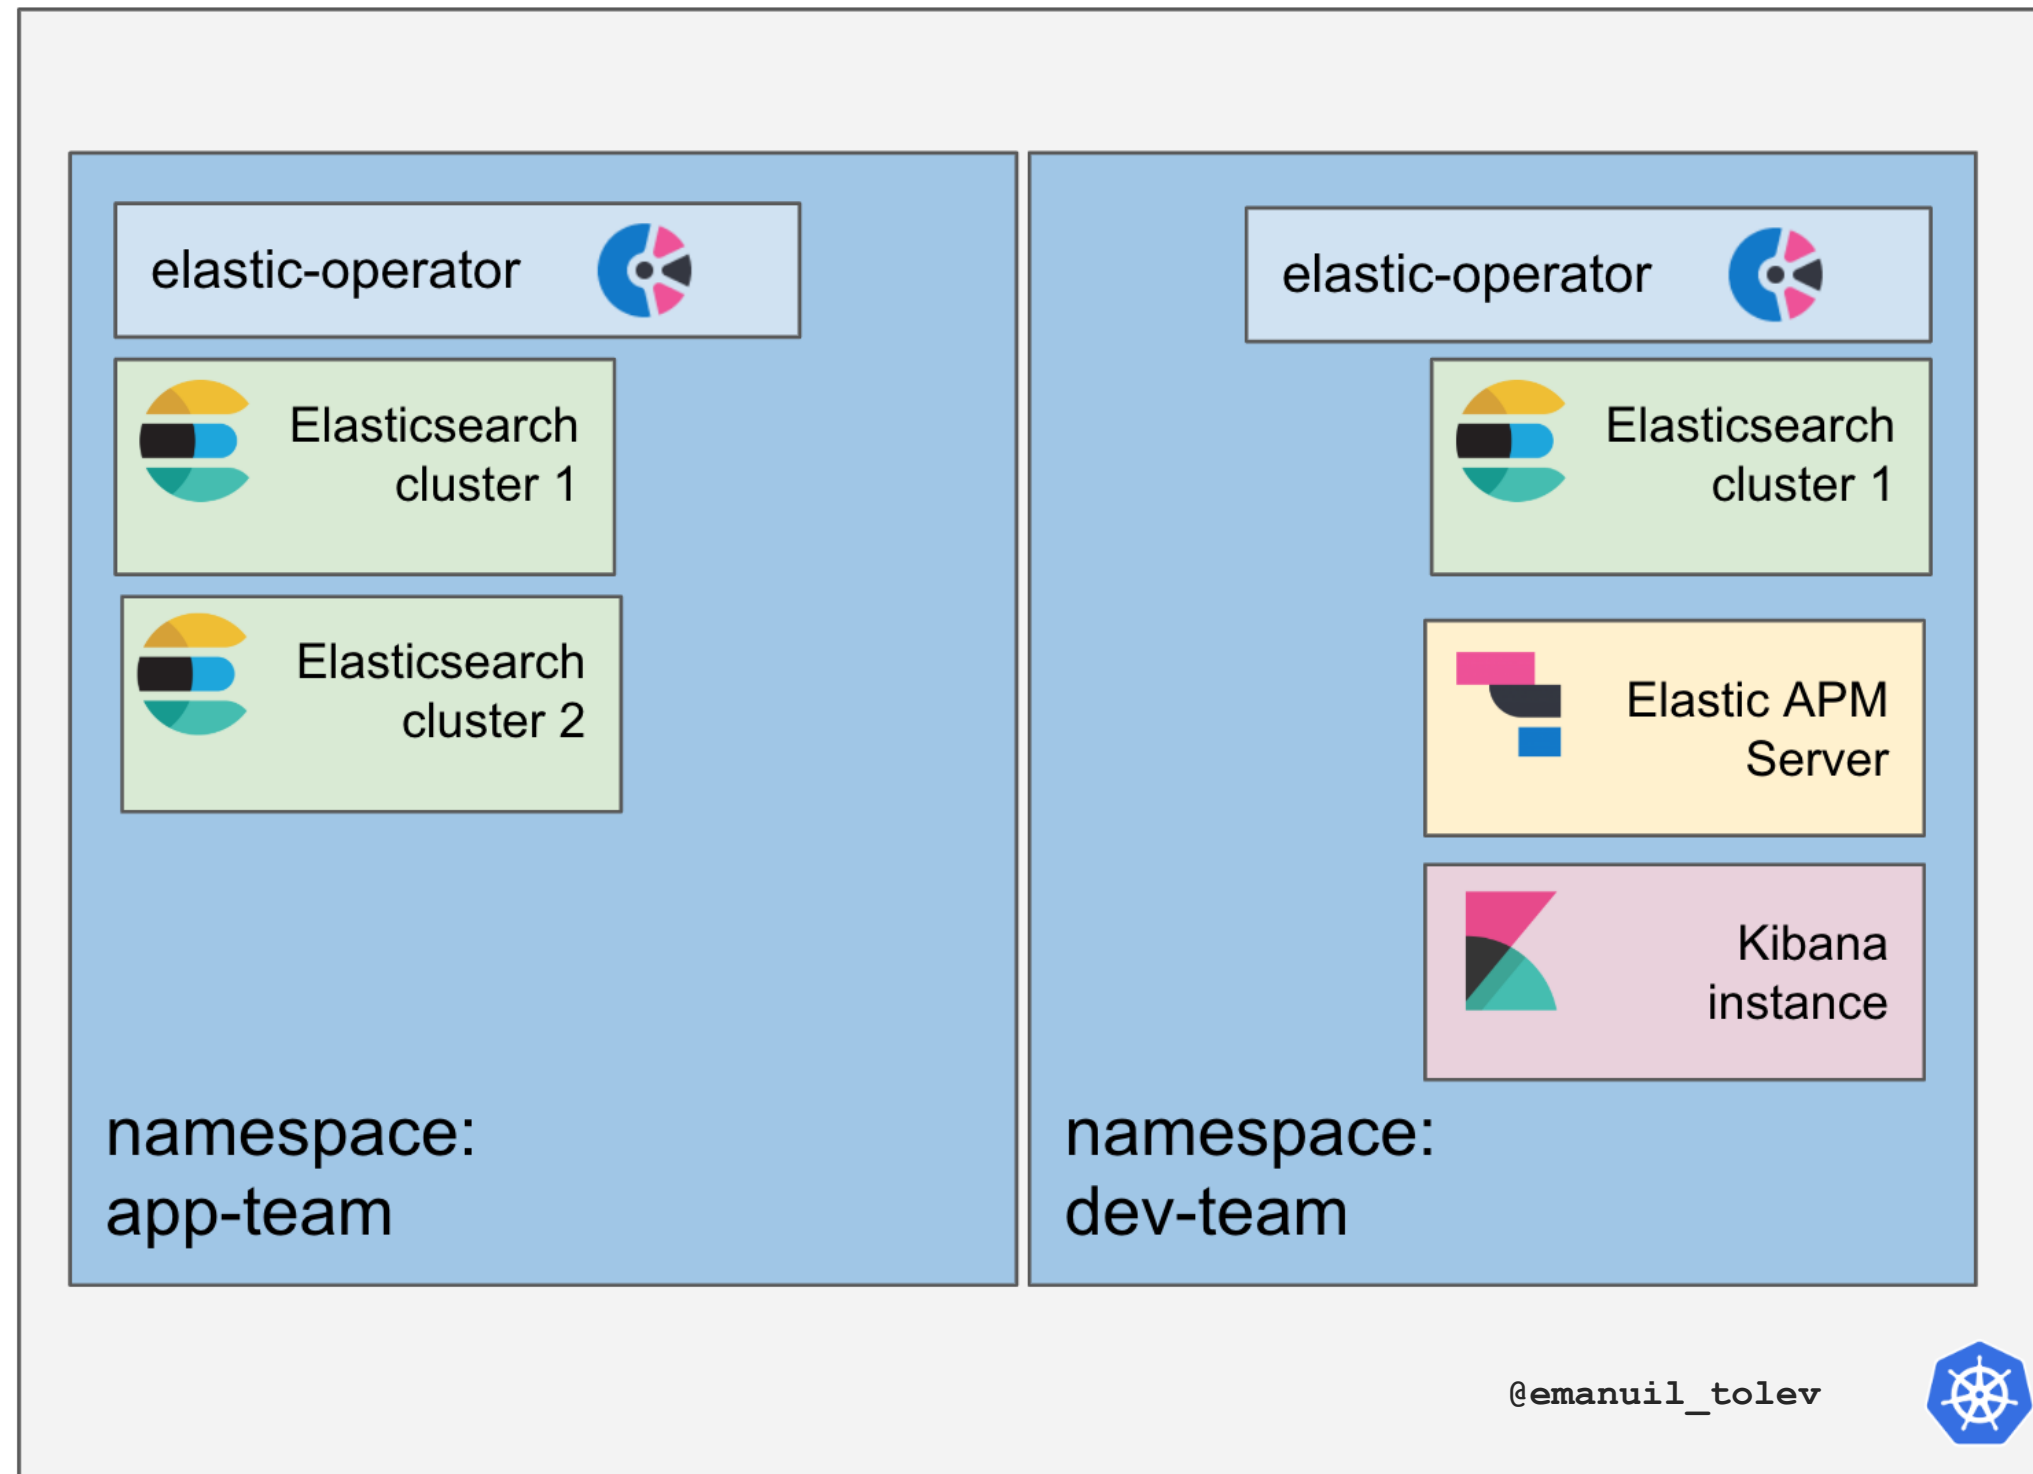
# Deployment

CRDs require cluster admin level permissions to  
install

# Global Namespace



# Single Namespace





Couchbase Operator  
provided by Couchbase

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



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 portable



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?"**

# Effective collaboration and solving production problems

Remember why you're doing all this

# Helm Charts vs Operator

# Where to next?

- Deeper look at the operator: <https://www.elastic.co/blog/introducing-elastic-cloud-on-kubernetes-the-elasticsearch-operator-and-beyond>
- The source code: <https://github.com/elastic/cloud-on-k8s>
- The slides: <https://noti.st/emanuil-tolev/CqvknF/from-containers-to-kubernetes-operators>

# Elastic interest?

community.elastic.co  
free lunch info sessions, ping me

# Questions & Discussion

@emanuil\_tolev  
etolev@elastic.co