



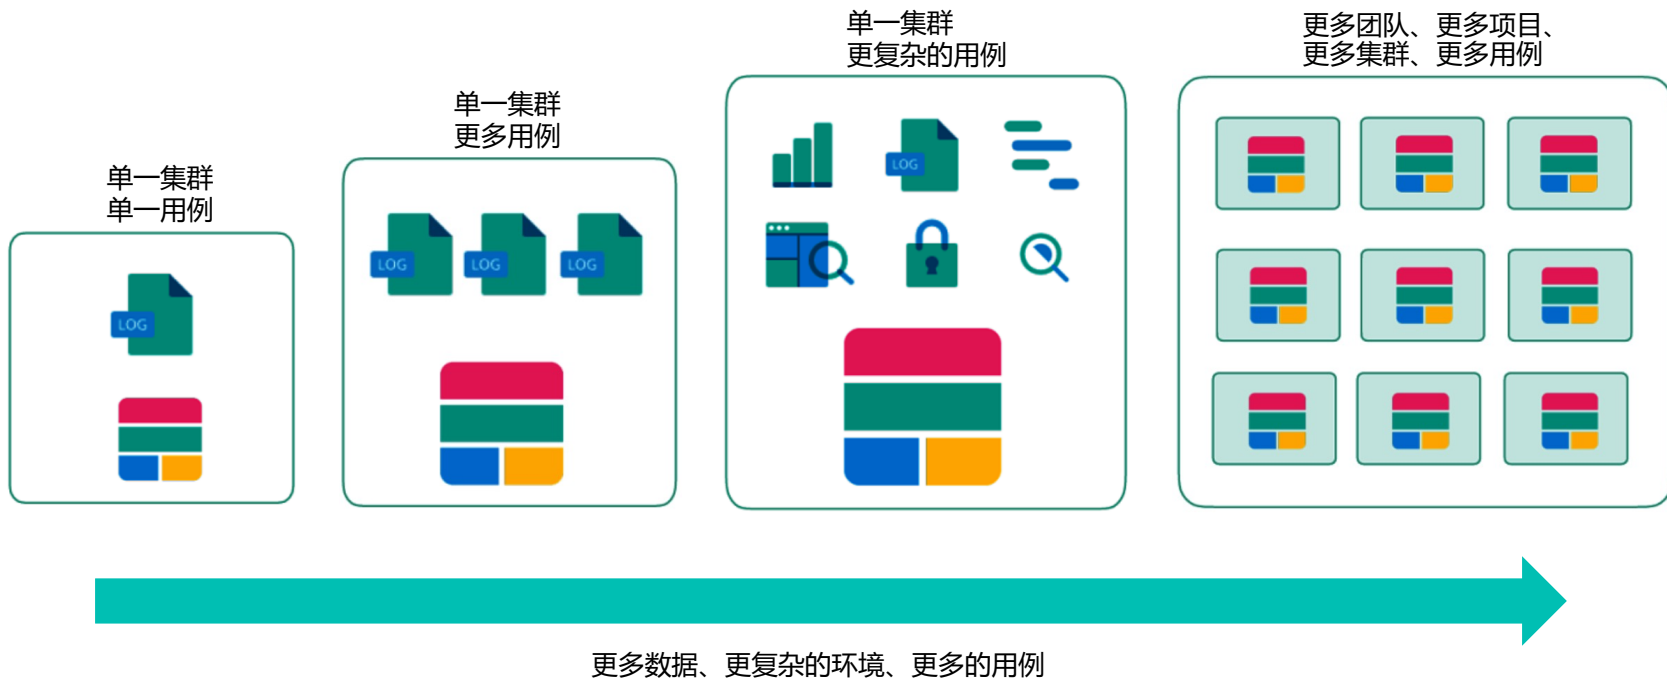
elastic 中国开发者大会 2023

# Elastic Stack在Kubernetes上的最佳实践分享

---

刘凡，架构师-容器/DevOps技术专家  
谷歌，2023/4/8

# 典型的采用Elastic Stack的过程



# Kubernetes是一个容器编排平台

✓ Kubernetes 是**基础平台**

✓ Kubernetes 是**可扩展**

✓ Kubernetes **技术成熟**

**Kubernetes...**

## 让每天的 工作更轻松



# Stateless is Easy, Stateful is Hard

# 在Kubernetes上运行ES的主要挑战



- 管理不同的**资源**：Statefulset , Services , Secrets
- **存储**：卷的管理
- 确保**一致性和可靠性**
- 扩容、缩容
- 更改**配置和升级**

# 深入了解Operator

# Kubernetes Operators

K8s Operator extends the k8s API with Custom Resource Definitions (CRD) that encapsulates operational domain knowledge and enable common day-N operations

**Operator = k8s API + CRD(declarative spec) + custom controller**

Suitable for stateful applications where the models provided by k8s are not sufficient to handle fault remediation, scaling with data integrity and availability.

Custom controller watches these Custom Resources and take actions to move the system to the desired state.



# Custom Resource

- Allows to introduce your own APIs in a cluster
- Custom Resource Definition (CRD) manifest
  - API (group, version, kind) and subresources
  - OpenAPI validation
  - Printer Column Support
  - Version Conversion
  - Defaulting
- Hosted REST endpoints
- Similar to built-in resources

```
apiVersion: db.example.com/v1
```

```
...
```

```
Spec:
```

```
  readReplicas: 2
```

```
  storage: fast
```

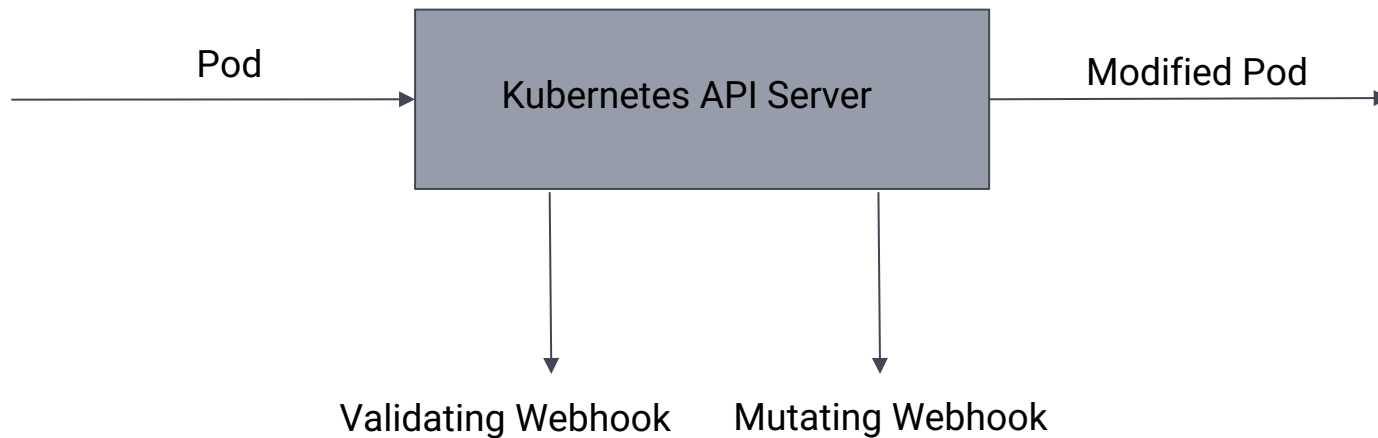
```
  Backup: weekly
```

```
  backupStorage: slow
```



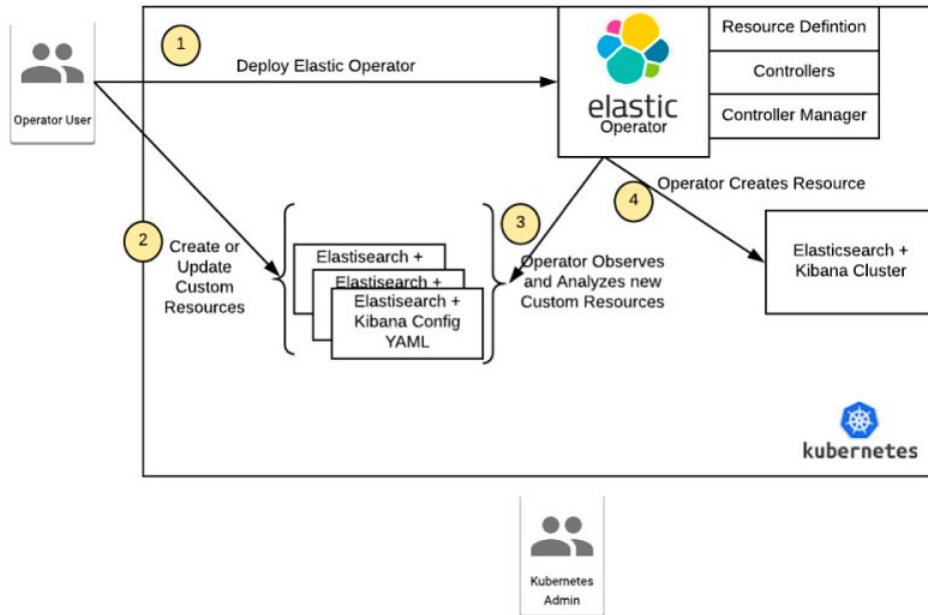
# Admission Webhooks

HTTP callbacks that receive admission requests

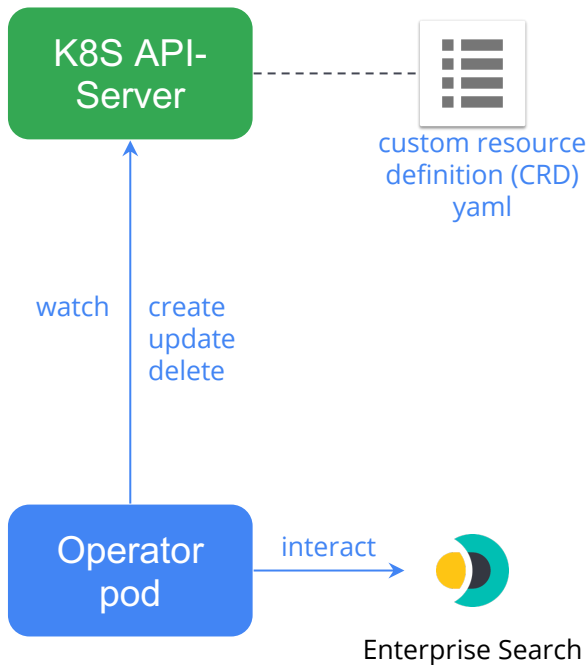


# ES Kubernetes Operator

- 一种交付K8S原生应用的方式
- 包含了运维的逻辑，比如知道如何进行ElasticSearch的升级



# Operator



## New event

A watched resource was created/updated/deleted



## Reconcile!

Get resource spec


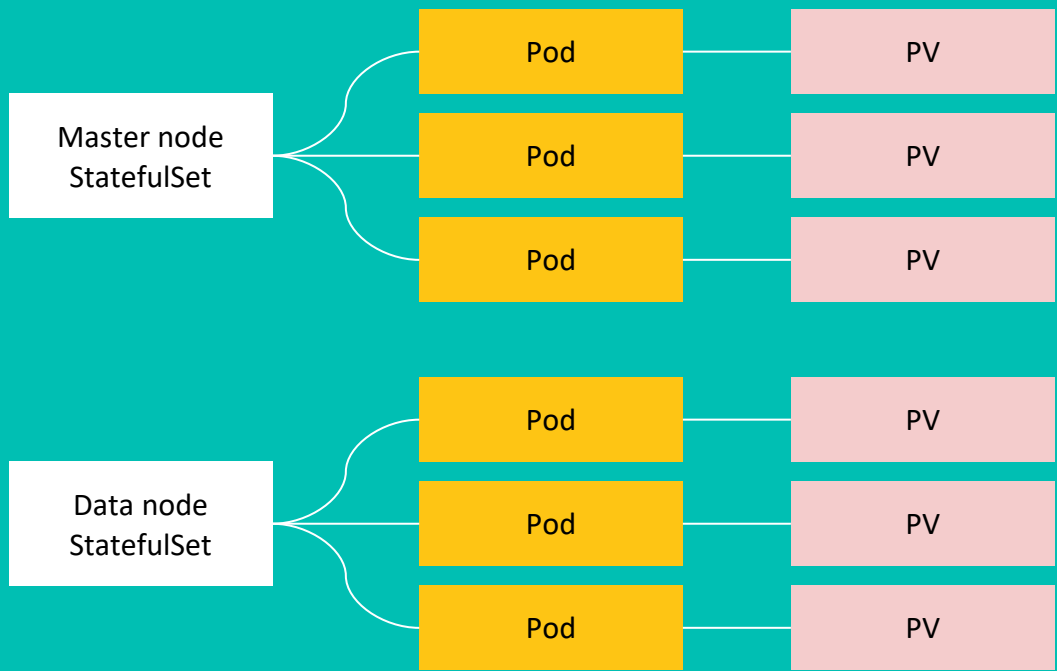
Create/update expected resources

- HTTP service
- TLS certificates (secret)
- Users (secret)
- Keystore (secret)
- ES configuration (configmap)
- ES Nodes StatefulSets (Pod + PersistentVolume)

Interact with the Elasticsearch API

- Change cluster settings
- Migrate data
- Prepare for rolling upgrade

```
1 apiVersion: elasticsearch.k8s.elastic.co/v1
2 kind: Elasticsearch
3 metadata:
4   name: es-sample
5 spec:
6   version: 7.15.2
7   nodeSets:
8   - name: zone-a-master
9     count: 3
10    config:
11      node.master: true
12      node.data: false
13    volumeClaimTemplates:
14    - metadata:
15      name: elasticsearch-data
16      spec:
17        accessModes:
18        - ReadWriteOnce
19        resources:
20          requests:
21            storage: 80Gi
22            storageClassName: pd-ssd
23    - name: zone-a-data
24      count: 3
25      config:
26        node.master: false
27        node.data: true
28      volumeClaimTemplates:
29      - metadata:
30        name: elasticsearch-data
31        spec:
32          accessModes:
33          - ReadWriteOnce
34          resources:
35            requests:
36              storage: 512Gi
37              storageClassName: pd-ssd
```



# Best practices for K8S Operator

If you're writing your own operator to manage a Kubernetes application, here are some best practices we recommend.

1. Develop one Operator per application
2. Use an SDK like [Kubebuilder](#)
3. Use declarative APIs
4. Compartmentalize features via multiple controllers
5. Use asynchronous sync loops

## Monitoring and logging your applications

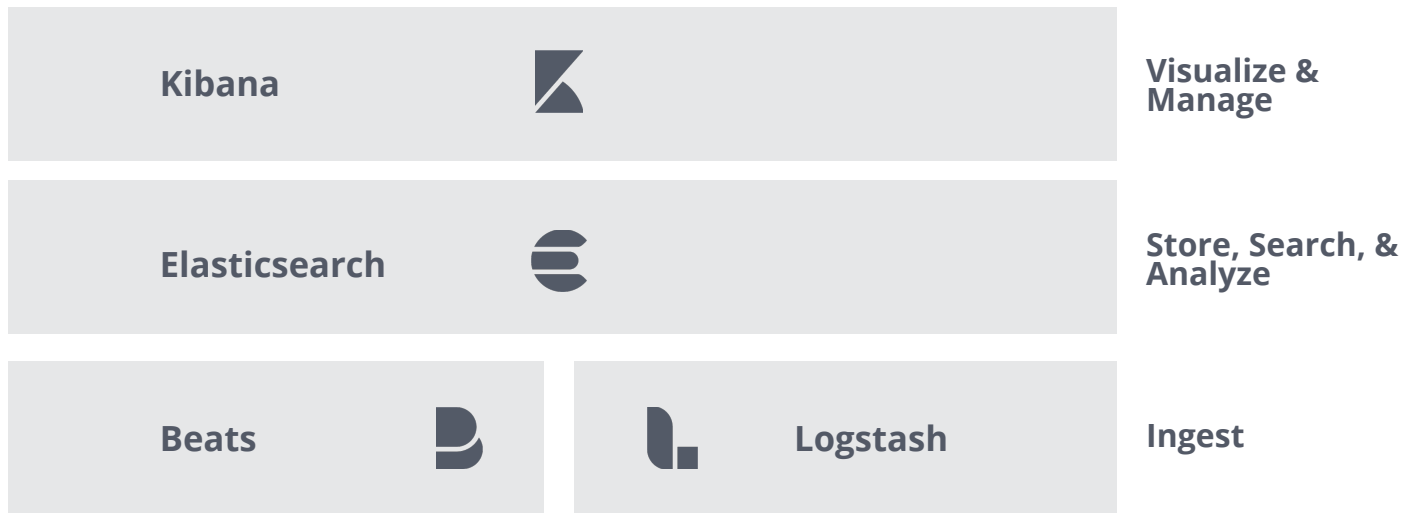
- Perform application-, node- and cluster-level log aggregation
- Properly label your metrics for easier view, aggregation, and analysis
- Expose application metrics via pod endpoints for scraping purposes

# What is Kubebuilder ?

The screenshot shows the GitHub repository page for `kubernetes-sigs / kubebuilder`. At the top right, there are buttons for `Watch` (51), `Star` (1,119), and `Fork` (203). Below these are navigation tabs for `Code` (selected), `Issues` (102), `Pull requests` (16), `Projects` (4), and `Insights`. The repository description is "Kubebuilder - SDK for building Kubernetes APIs using CRDs" with a link to <http://book.kubebuilder.io>. A tag `k8s-sig-api-machinery` is visible. At the bottom, a bar displays statistics: `721 commits`, `3 branches`, `25 releases`, `60 contributors`, and `Apache-2.0` license.

# Elastic Stack 的部署选项

SOLUTIONS





















# OSS Elastic on K8S

elastic / cloud-on-k8s Public

<> Code Issues 302 Pull requests 12 Actions Security Insights

master 16 branches 25 tags

Go to file Add file Code

 renovate	Update module go.elastic.co/apm/module/apmelasticsearch to v1.15...	1343283 2 days ago	3,159 commits
 .ci	Update golang Docker tag to v1.17.5 (#5149)		2 days ago
 .github	Control ignorePaths of renovate config to include test/ (#4386)		9 months ago
 cmd	Exit main go routine even if there's no error (#5106)		9 days ago
 config	Stop using deprecated xpack.monitoring.* settings (#5136)		5 days ago
 deploy	Allow K8S node labels to be propagated as pod annotations (#5054)		12 days ago
 docs	Update module go.elastic.co/apm/module/apmelasticsearch to v1.15...		2 days ago
 hack	Update module helm.sh/helm/v3 to v3.7.2 (#5146)		2 days ago
 pkg	Stop using deprecated xpack.monitoring.* settings (#5136)		5 days ago
 support	Update golangci-lint/goimports configuration (#4945)		2 months ago
 test/e2e	Update golang Docker tag to v1.17.5 (#5149)		2 days ago
 .gitignore	Remove legacy v1beta1 CRDs (#5029)		24 days ago
 .golangci.yml	Update golangci-lint/goimports configuration (#4945)		2 months ago
 CONTRIBUTING.md	Update golangci-lint/goimports configuration (#4945)		2 months ago
 Dockerfile	Update golang Docker tag to v1.17.5 (#5149)		2 days ago
 LICENSE.txt	Elastic License 2.0 (#4887)		2 months ago
 Makefile	Bump Stack version to 7.15.2 (#5076)		11 days ago
 NOTICE.txt	Update module go.elastic.co/apm/module/apmelasticsearch to v1.15...		2 days ago










# OSS Elastic on GKE

elasticsearch-cn / elastic-on-gke Public

<> Code This path skips through empty directories Actions Projects Wiki Security Insights

develop 2 branches 2 tags

Go to file Add file Code

 bindiego	added terraform sample for provisioning the GKE cluster	78c3e2a 11 days ago	35 commits
 bin	added terraform sample for provisioning the GKE cluster		11 days ago
 conf	ECK/ELK/GKE upgrades		3 months ago
 dingoperf	migrated from <a href="https://github.com/bindiego/local_services/tree/develop">https://github.com/bindiego/local_services/tree/develop</a>		13 months ago
 templates	upgraded both Elastic Stack version to 7.15.2 & GKE to the latest sta...		17 days ago
 terraform/gcp	added terraform sample for provisioning the GKE cluster		11 days ago
 .gitignore	migrated from <a href="https://github.com/bindiego/local_services/tree/develop">https://github.com/bindiego/local_services/tree/develop</a>		13 months ago
 Makefile	added support APM Server		7 months ago
 readme.md	minor updates in readme		17 days ago

readme.md

This project is migrated from:

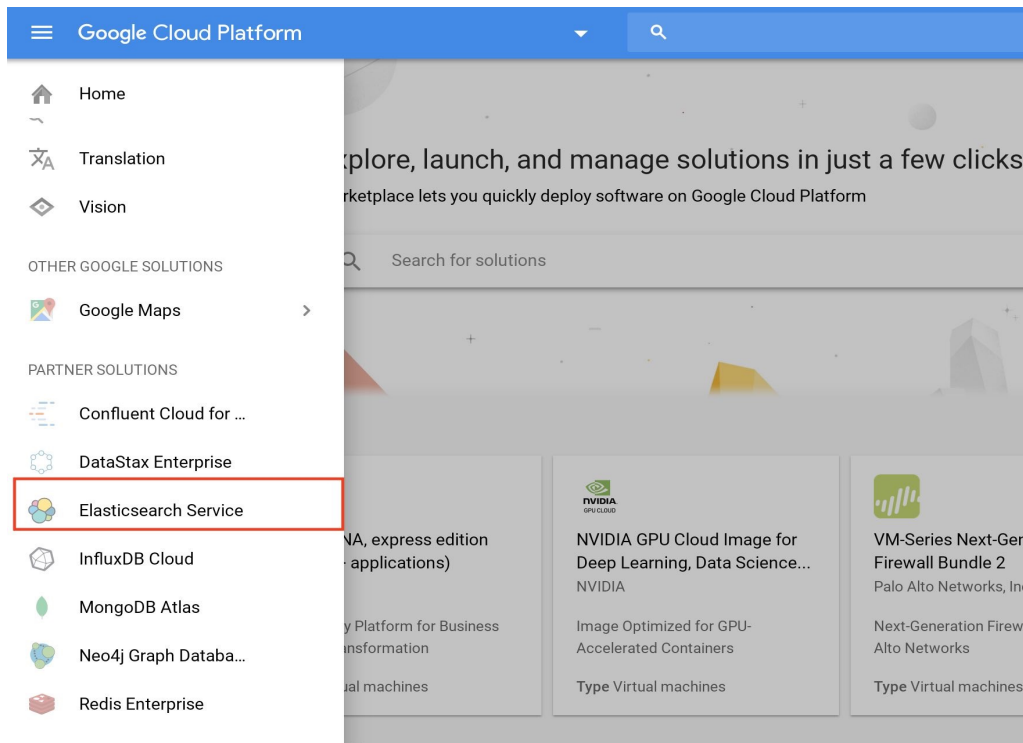
[https://github.com/bindiego/local\\_services/tree/develop/k8s/gke/elastic](https://github.com/bindiego/local_services/tree/develop/k8s/gke/elastic)

## Elastic Stack on k8s/GKE

This project is currently in a MVP (Minimum Viable Product) status. So the entire process may involve some manual setups (find & replace in a text editor). Once you familiar with the structure, you will find it's extremely easy. I am not sure if [Helm](#) can achieve the flexibility for this kind of data layer products. Let's start with simplicity anyway.



# Elastic Cloud is accessible directly from the Google Cloud Console



# Centralized Observability for Hybrid/Multi Cloud


Monitoring Anthos GKE

# Anthos is Google's platform for application modernization

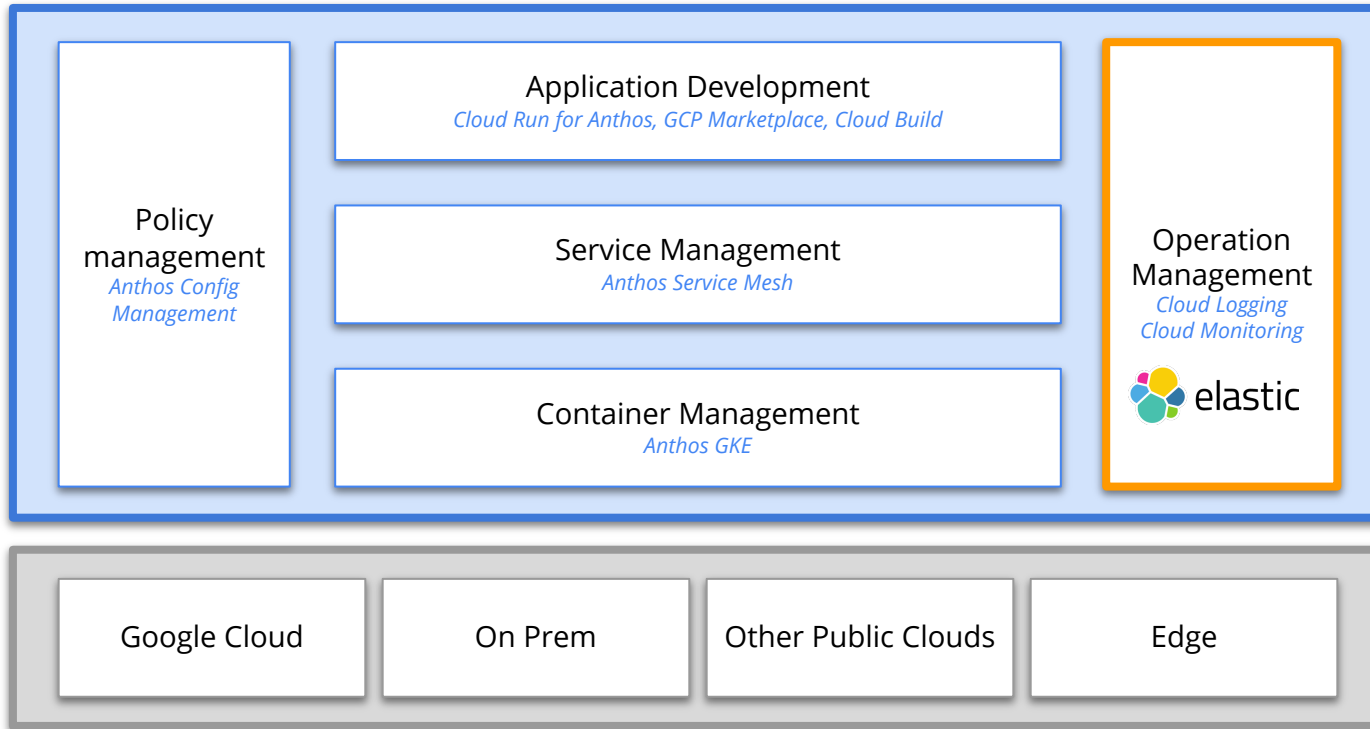


“ Anthos' multi-cloud capabilities are so far unique in the industry

Forbes, April 2019



# Core components of Anthos



# Elastic has launched new observability features for Google cloud that make this pattern easier to use

Google Cloud module - now in beta

- **Metricbeat**
  - Compute, Load balancing, Pub/Sub, Operations and Storage Metricsets
- **Filebeat**
  - Audit, VPC flow, and Firewall logs exported from Operations to a Pub/Sub sink



**THANK YOU**

---





专业、垂直、纯粹的 Elastic 开源技术交流社区

<https://elasticsearch.cn/>