QINGCLOUD | KUBESPHERE

Effective Logging in Multi-Tenant Kubernetes Environment

Benjamin Huo
Lead of KubeSphere Observability team

Dan Ma
Senior SE of KubeSphere Observability team

QINGCLOUD | KUBESPHERE

How to debug your K8S workloads?

✓kubectl logs pod1 -c container2

✓Cloud providers' logging solutions: stackdriver, cloudwatch

✓ISVs' logging solution: Splunk, Sumo Logic, Datadog

✓OSS logging solutions: EFK, Loki

QINGCLOUD | KUBESPHERE

How to debug your K8S workloads?

	Pros	Cons
Stackdriver/Cloudwatch logs	Easy to use Low maintainence efforts	Vendor lock-in High cost Public cloud only
Splunk/Sumo Logic/Datadog	Easy to use Low maintainence efforts	Vendor lock-in High cost
EFK	The most popular and mature OSS logging solution	High resource footprint
Grafana loki	Emerging star Low cost	Slower log content search Need time to become mature

QINGCLOUD | KUBESPHERE

K8S common logging architecture

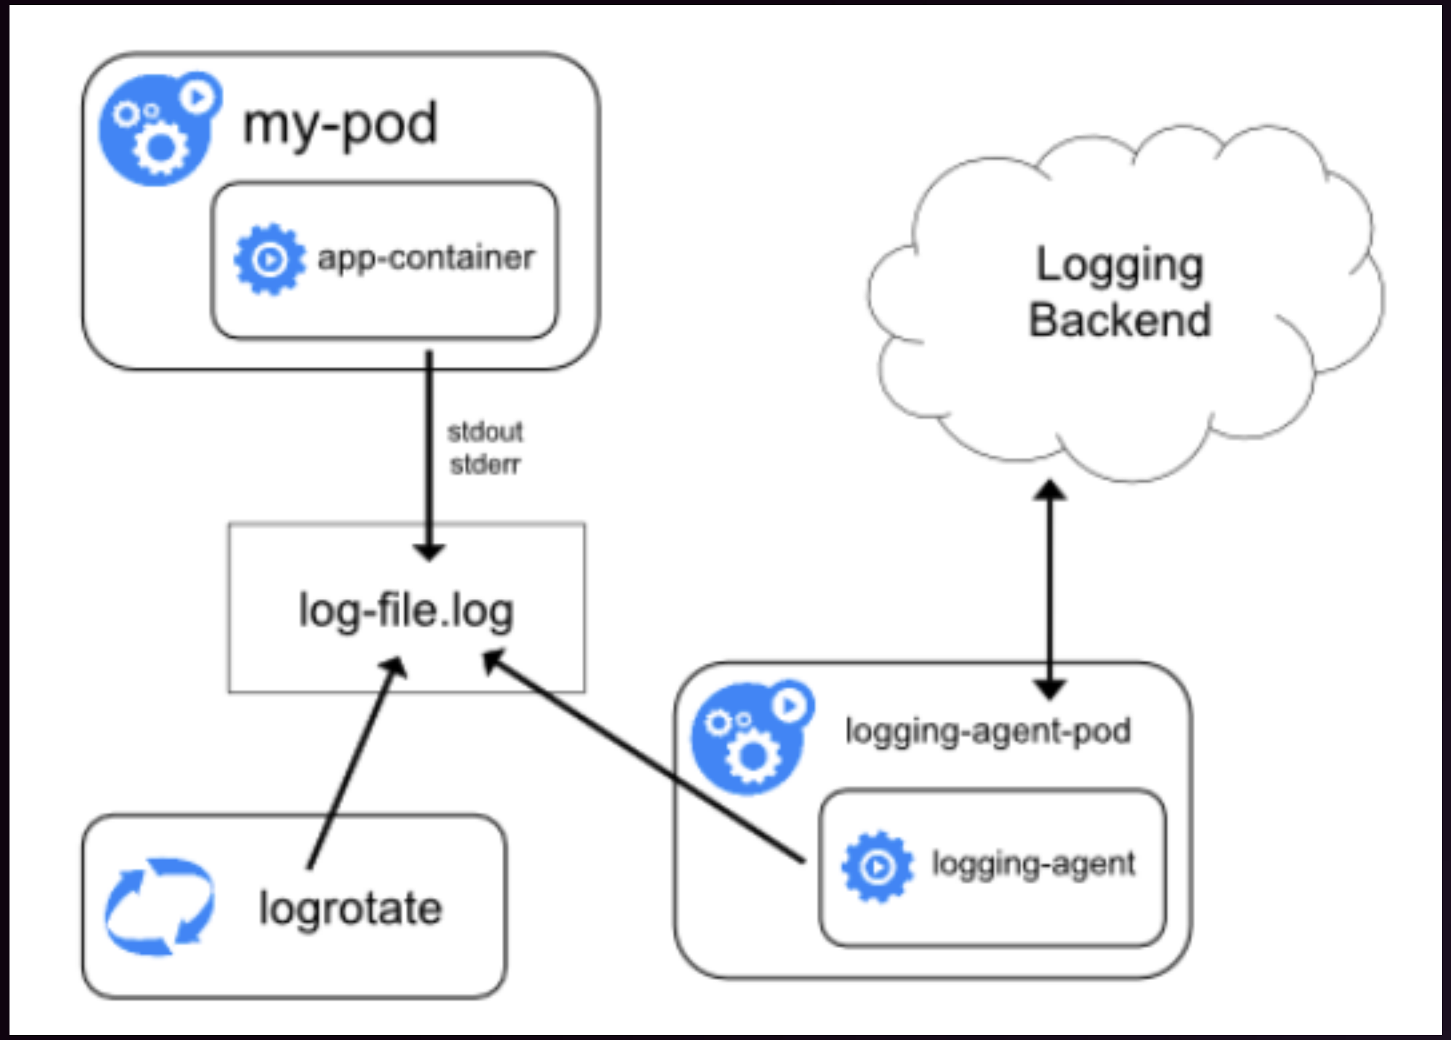

- Logging Agent
- Logging Backend
- Logging Console

K8S distribution’s logging requirements

CNCF: A K8S distribution can be installed to
public(or private) cloud or bare mental

Only OSS logging solutions are qualified:

✓EFK

✓Loki

QINGCLOUD | KUBESPHERE

Enterprises' logging requirements

- ✓ Maturity
- ✓ Access control
- ✓ Ability to integrate with existing logging system
- ✓ Ability to integrate with big data platform
- ✓ K8S native

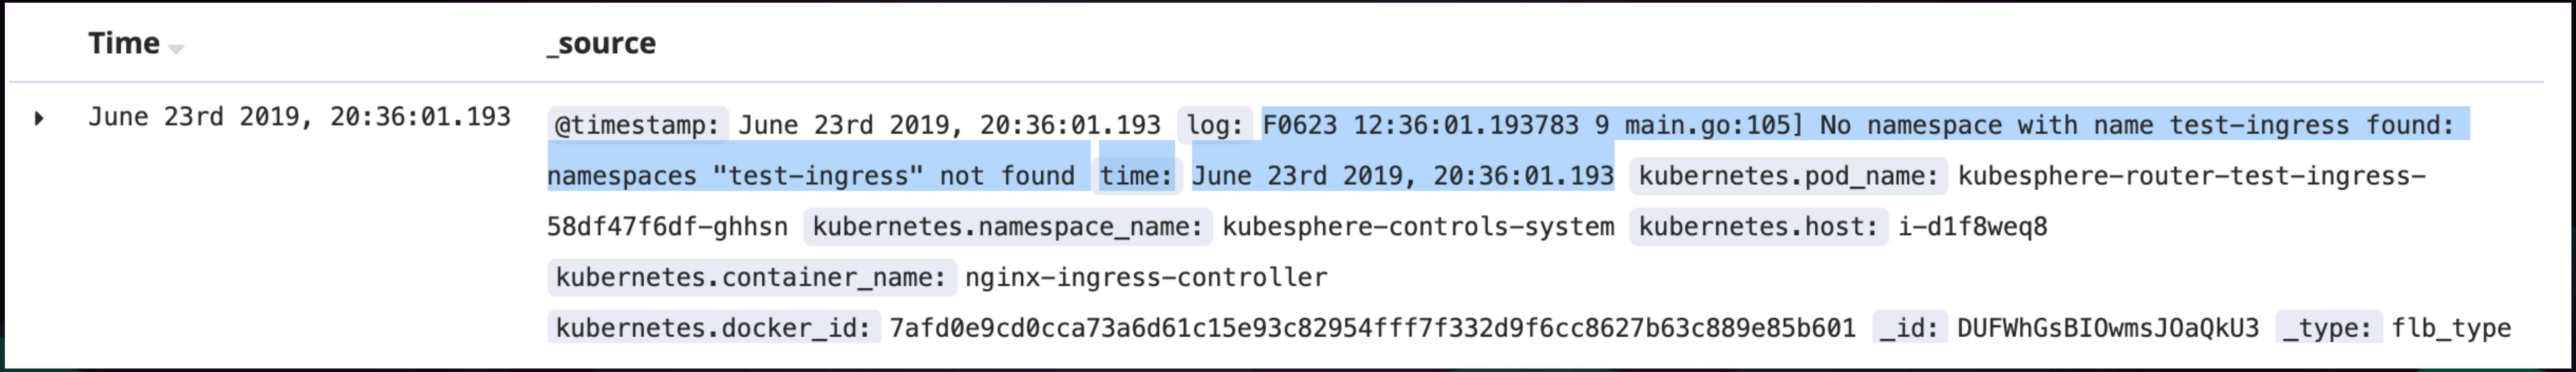

```
Time        _source
June 23rd 2019, 20:36:01.193   @timestamp: June 23rd 2019, 20:36:01.193 log: F0623 12:36:01.193783 9 main.go:105] No namespace with name test-ingress found:
namespaces "test-ingress" not found time: June 23rd 2019, 20:36:01.193 kubernetes.pod_name: kubesphere-router-test-ingress-
58df47f6df-ghhsn kubernetes.namespace_name: kubesphere-controls-system kubernetes.host: i-d1f8weq8
kubernetes.container_name: nginx-ingress-controller
kubernetes.docker_id: 7afd0e9cd0cca73a6d61c15e93c82954fff7f332d9f6cc8627b63c889e85b601 _id: DUFWhGsBIOwmsJOaQkU3 _type: flb_type
```

OSS logging solution analysis

	Maturity	Access Control	Integration	K8S native
EFK	High	Medium	Good	Poor
Loki	Low	Good	Poor	Good

QINGCLOUD | KUBESPHERE

K8S distribution + Enterprise users + Logging = ?

QINGCLOUD | KUBESPHERE

Logging solution of KubeSphere

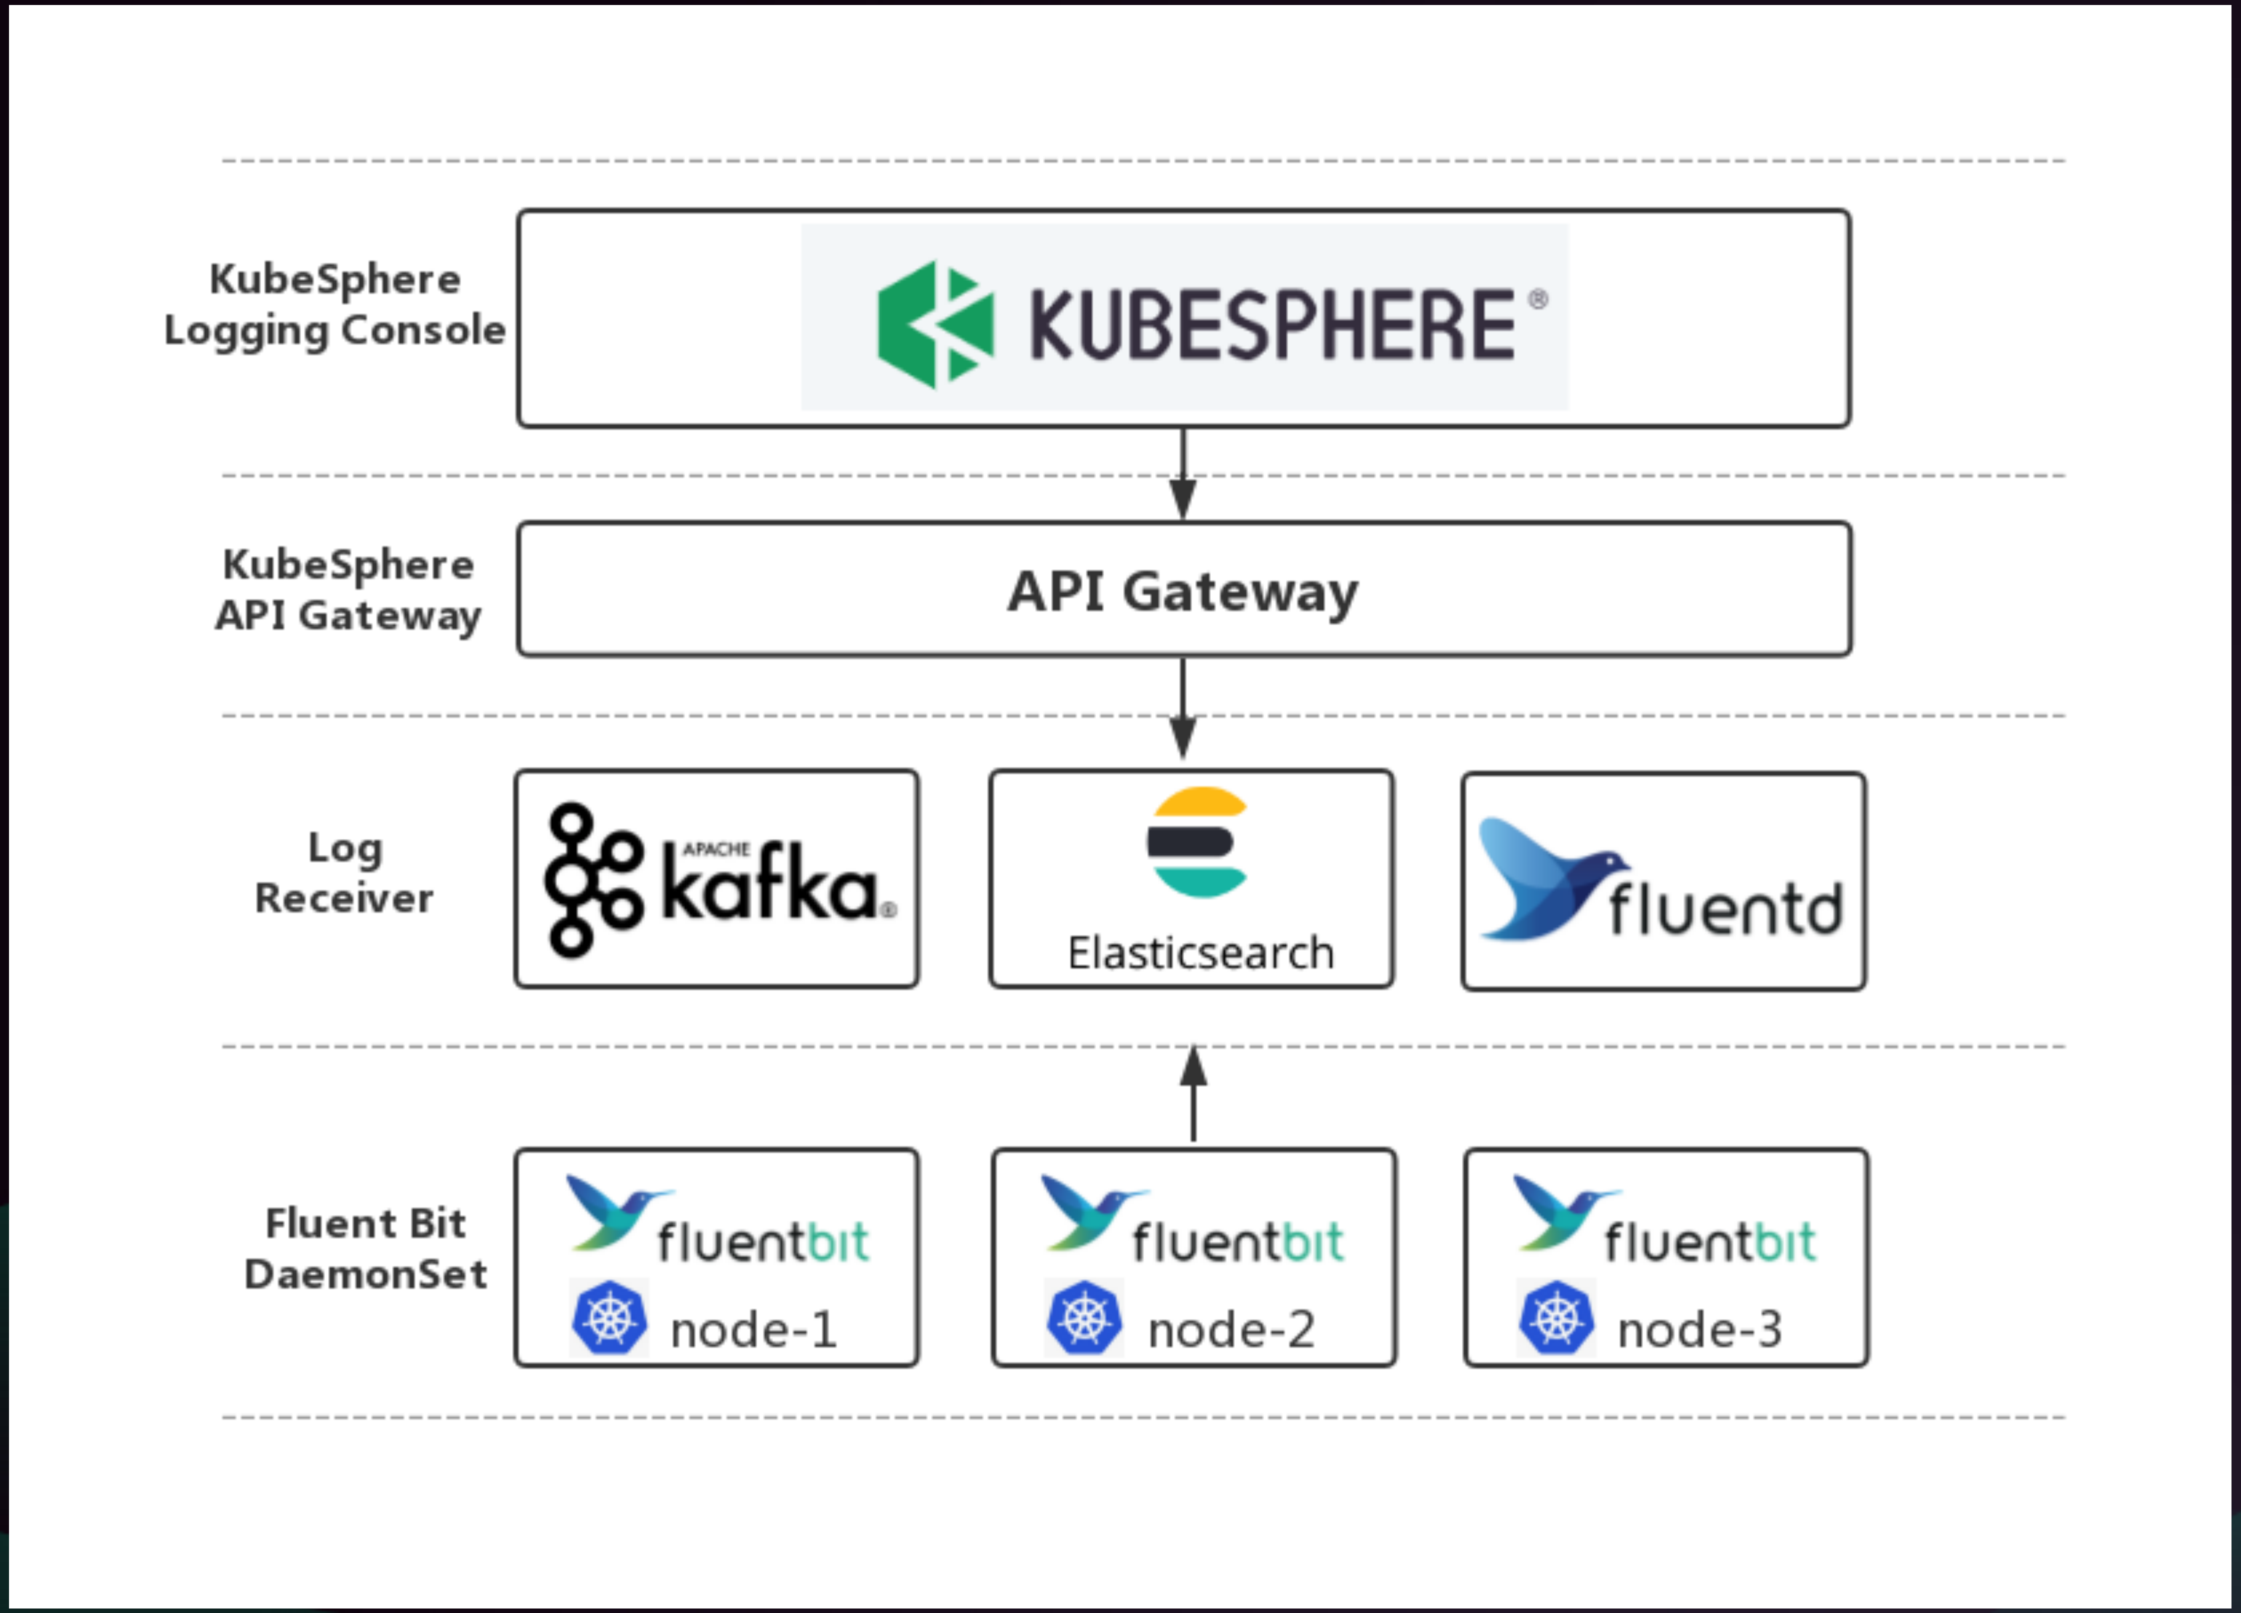

Agent: Fluent Bit instead of Fluentd

Backend: Elasticsearch

Console: Customized logging console

Integration: Kafka, Fluentd

K8S native: Search by ns, workload…

QINGCLOUD | KUBESPHERE

Multi-Tenant architecture

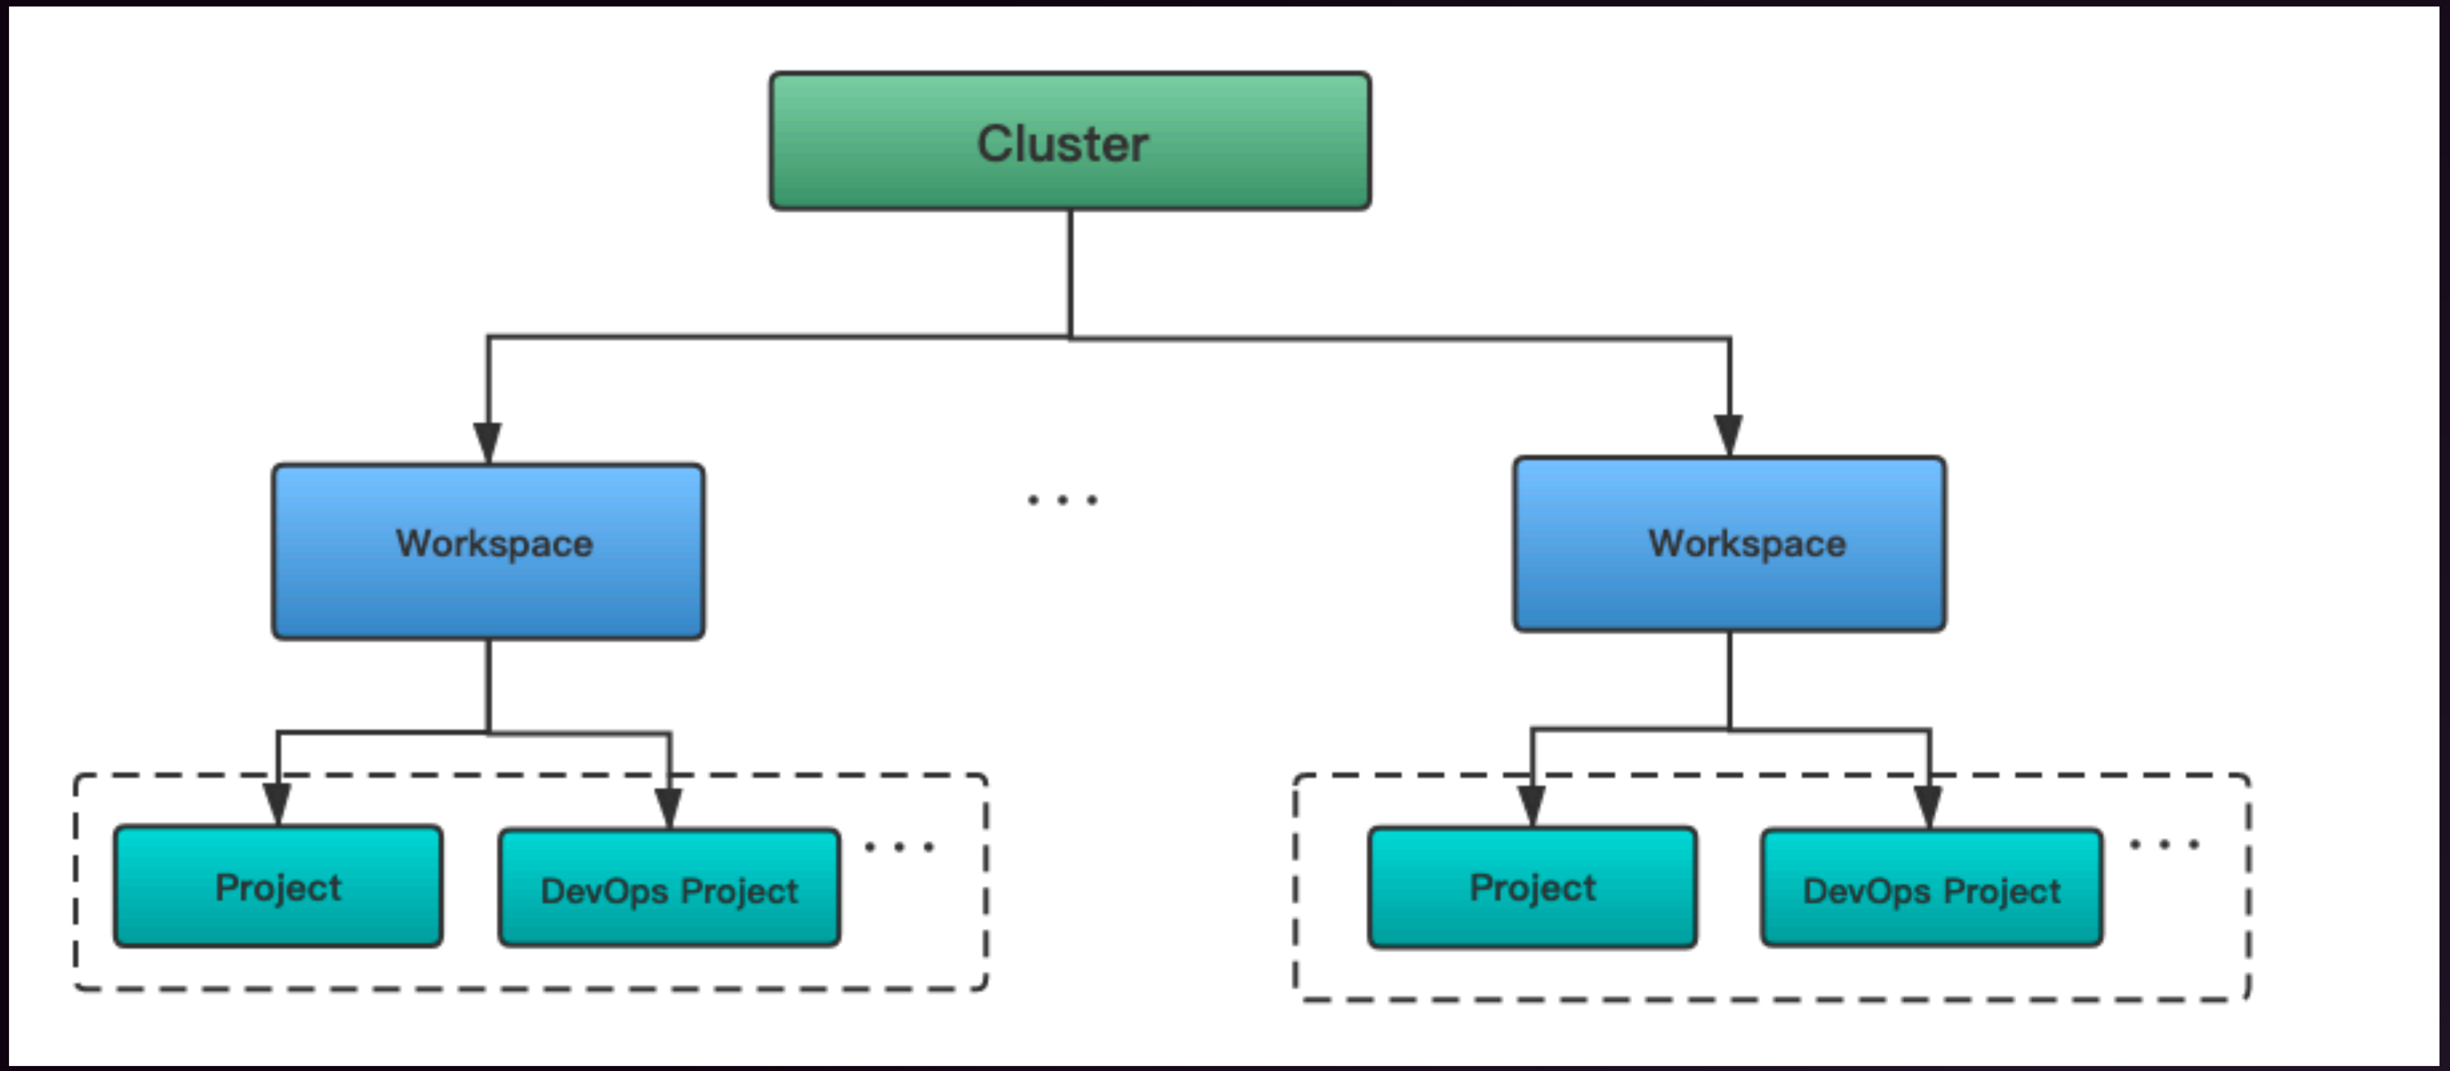

Cluster: admin/regular/workspace manager

Workspace: admin/regular/viewer

Project: admin/operator/viewer

3-Tier RBAC based multi-tenant architecture

QINGCLOUD | KUBESPHERE

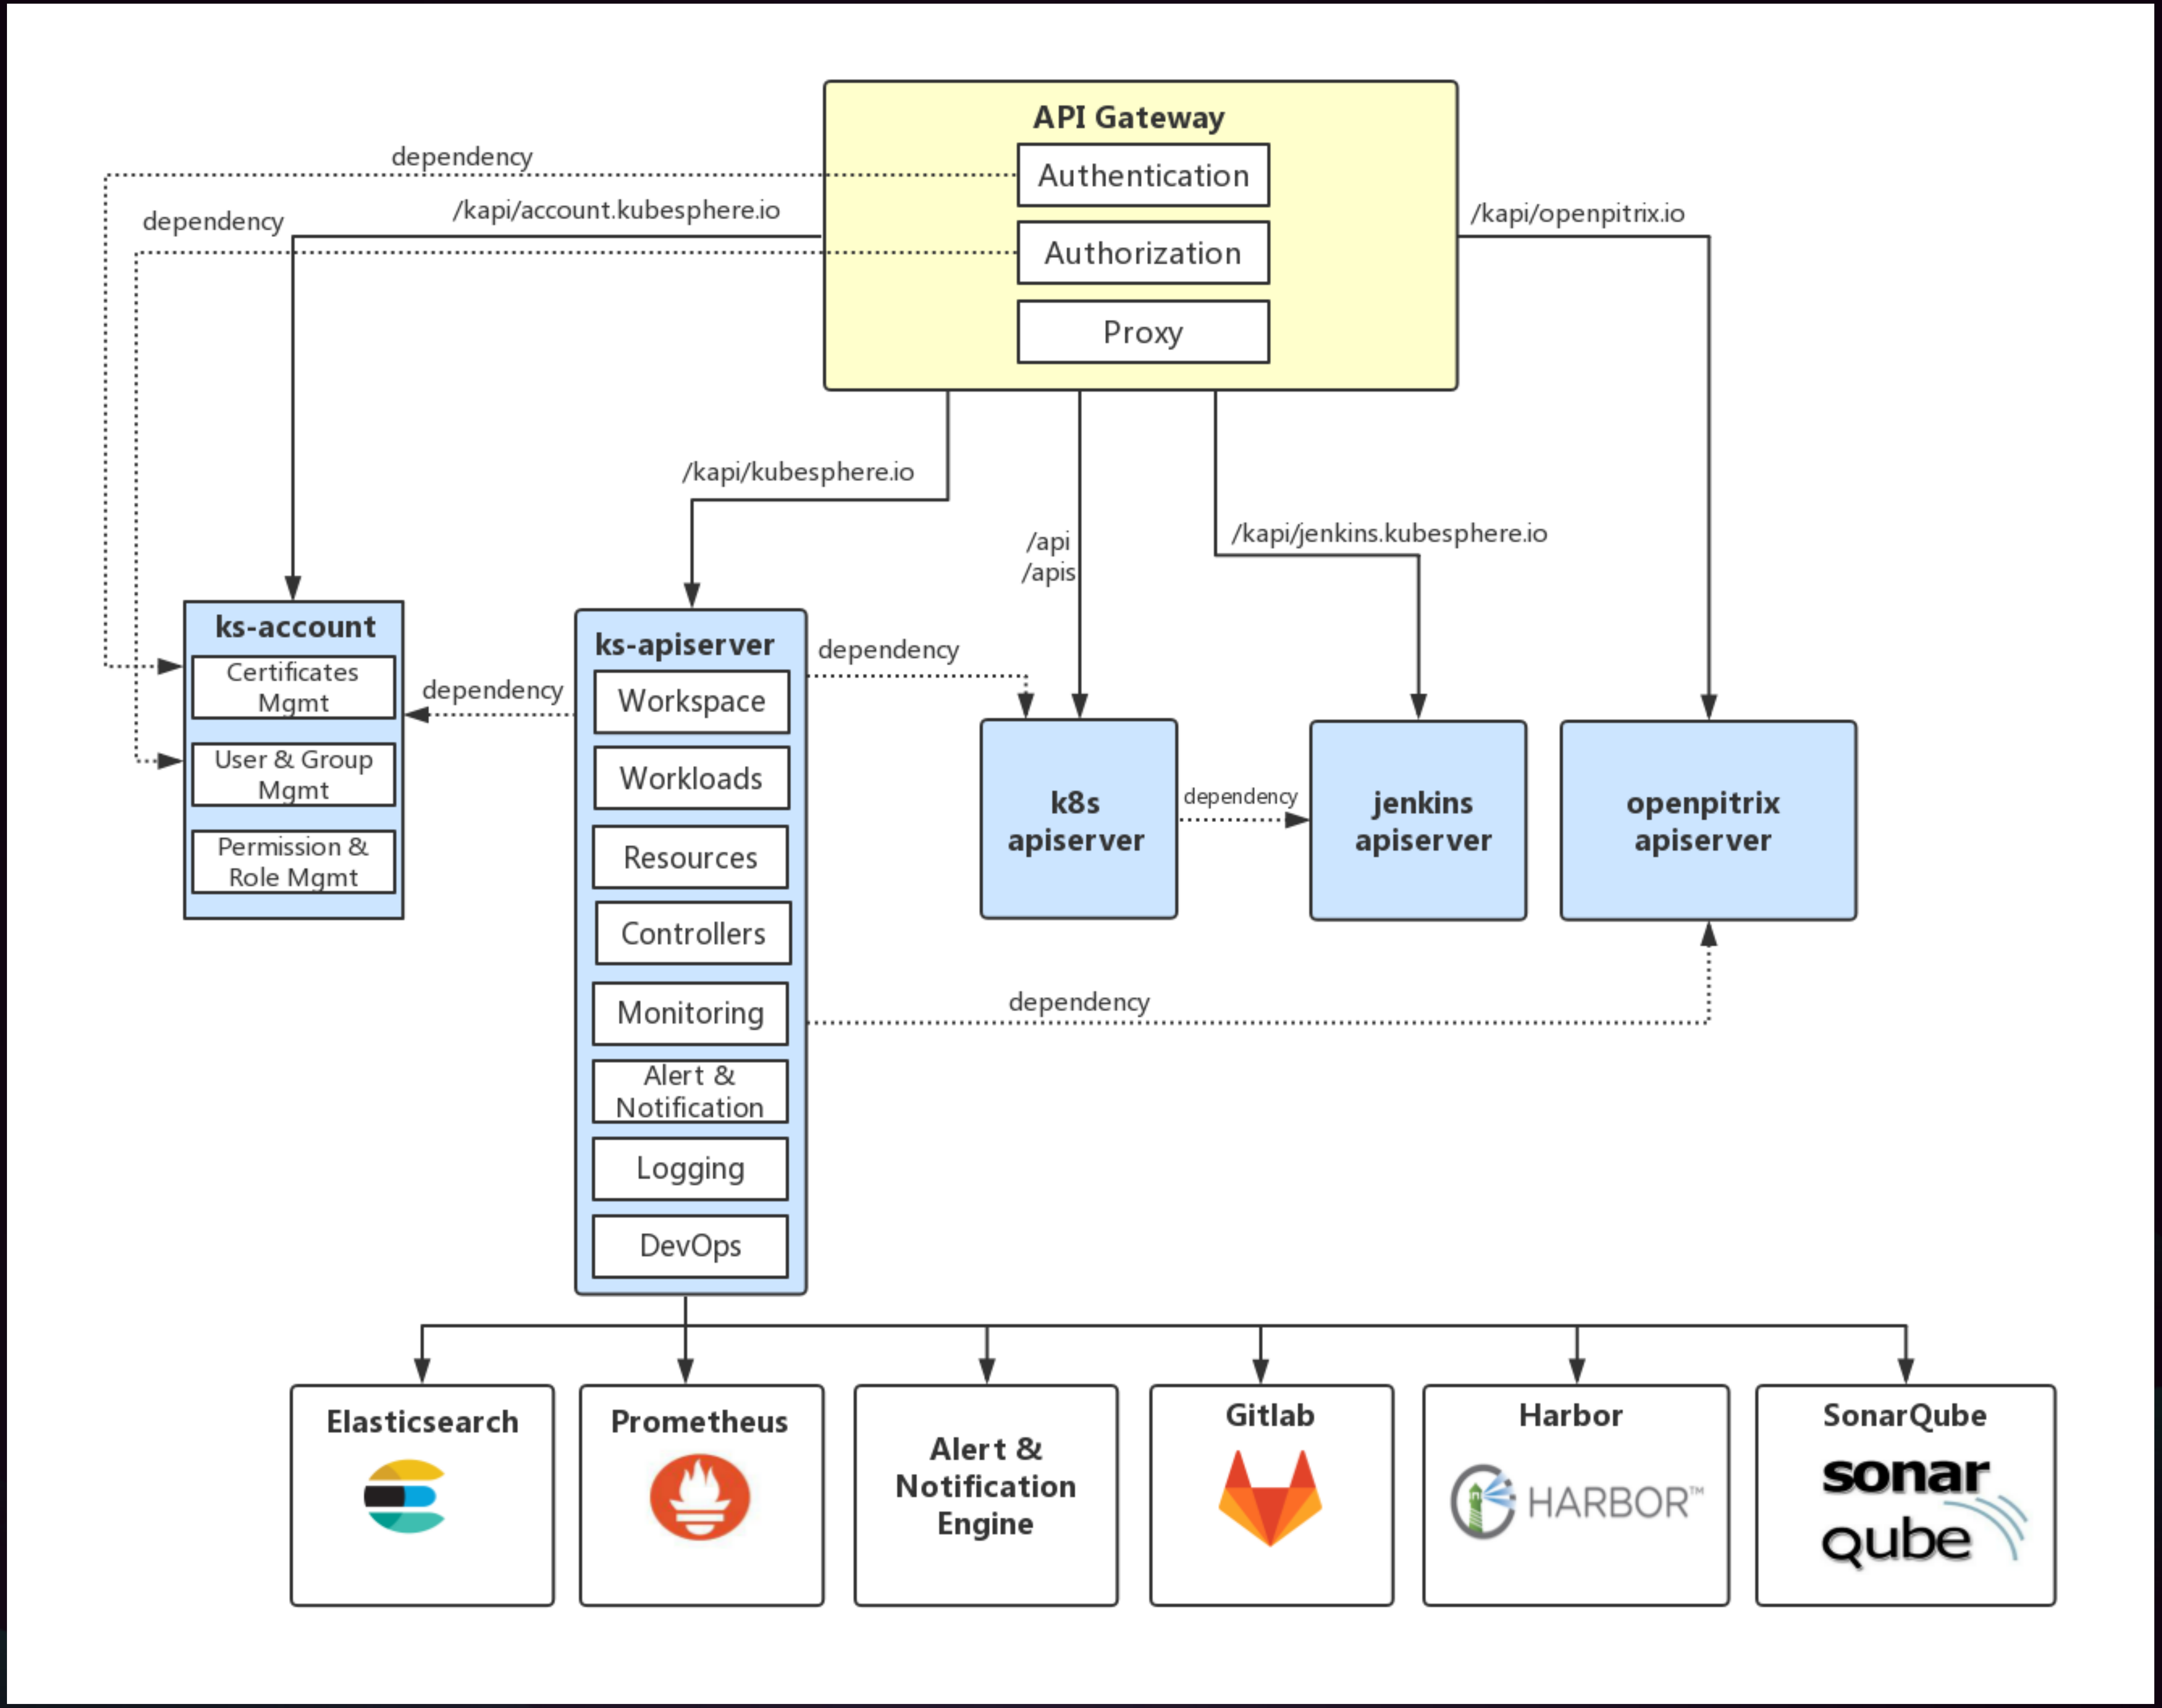

QINGCLOUD | KUBESPHERE

Multi-Tenant logging APIs

/kapis/logging.kubesphere.io/v1alpha2/cluster

/kapis/logging.kubesphere.io/v1alpha2/workspaces/{workspace}

/kapis/logging.kubesphere.io/v1alpha2/namespaces/{namespace}

/kapis/logging.kubesphere.io/v1alpha2/namespaces/{namespace}/workloads/{workload}

/kapis/logging.kubesphere.io/v1alpha2/namespaces/{namespace}/pods/{pod}

/kapis/logging.kubesphere.io/v1alpha2/namespaces/{namespace}/pods/{pod}/containers/{container}

Multi-Tenant benefits for logging

- ✓ User can only access logs he is authorized to
- ✓ Cluster admin manages log setting like:
  - Where to send the logs
  - Which logs should be collected
  - When to send the logs

QINGCLOUD | KUBESPHERE

Why Fluent Bit?

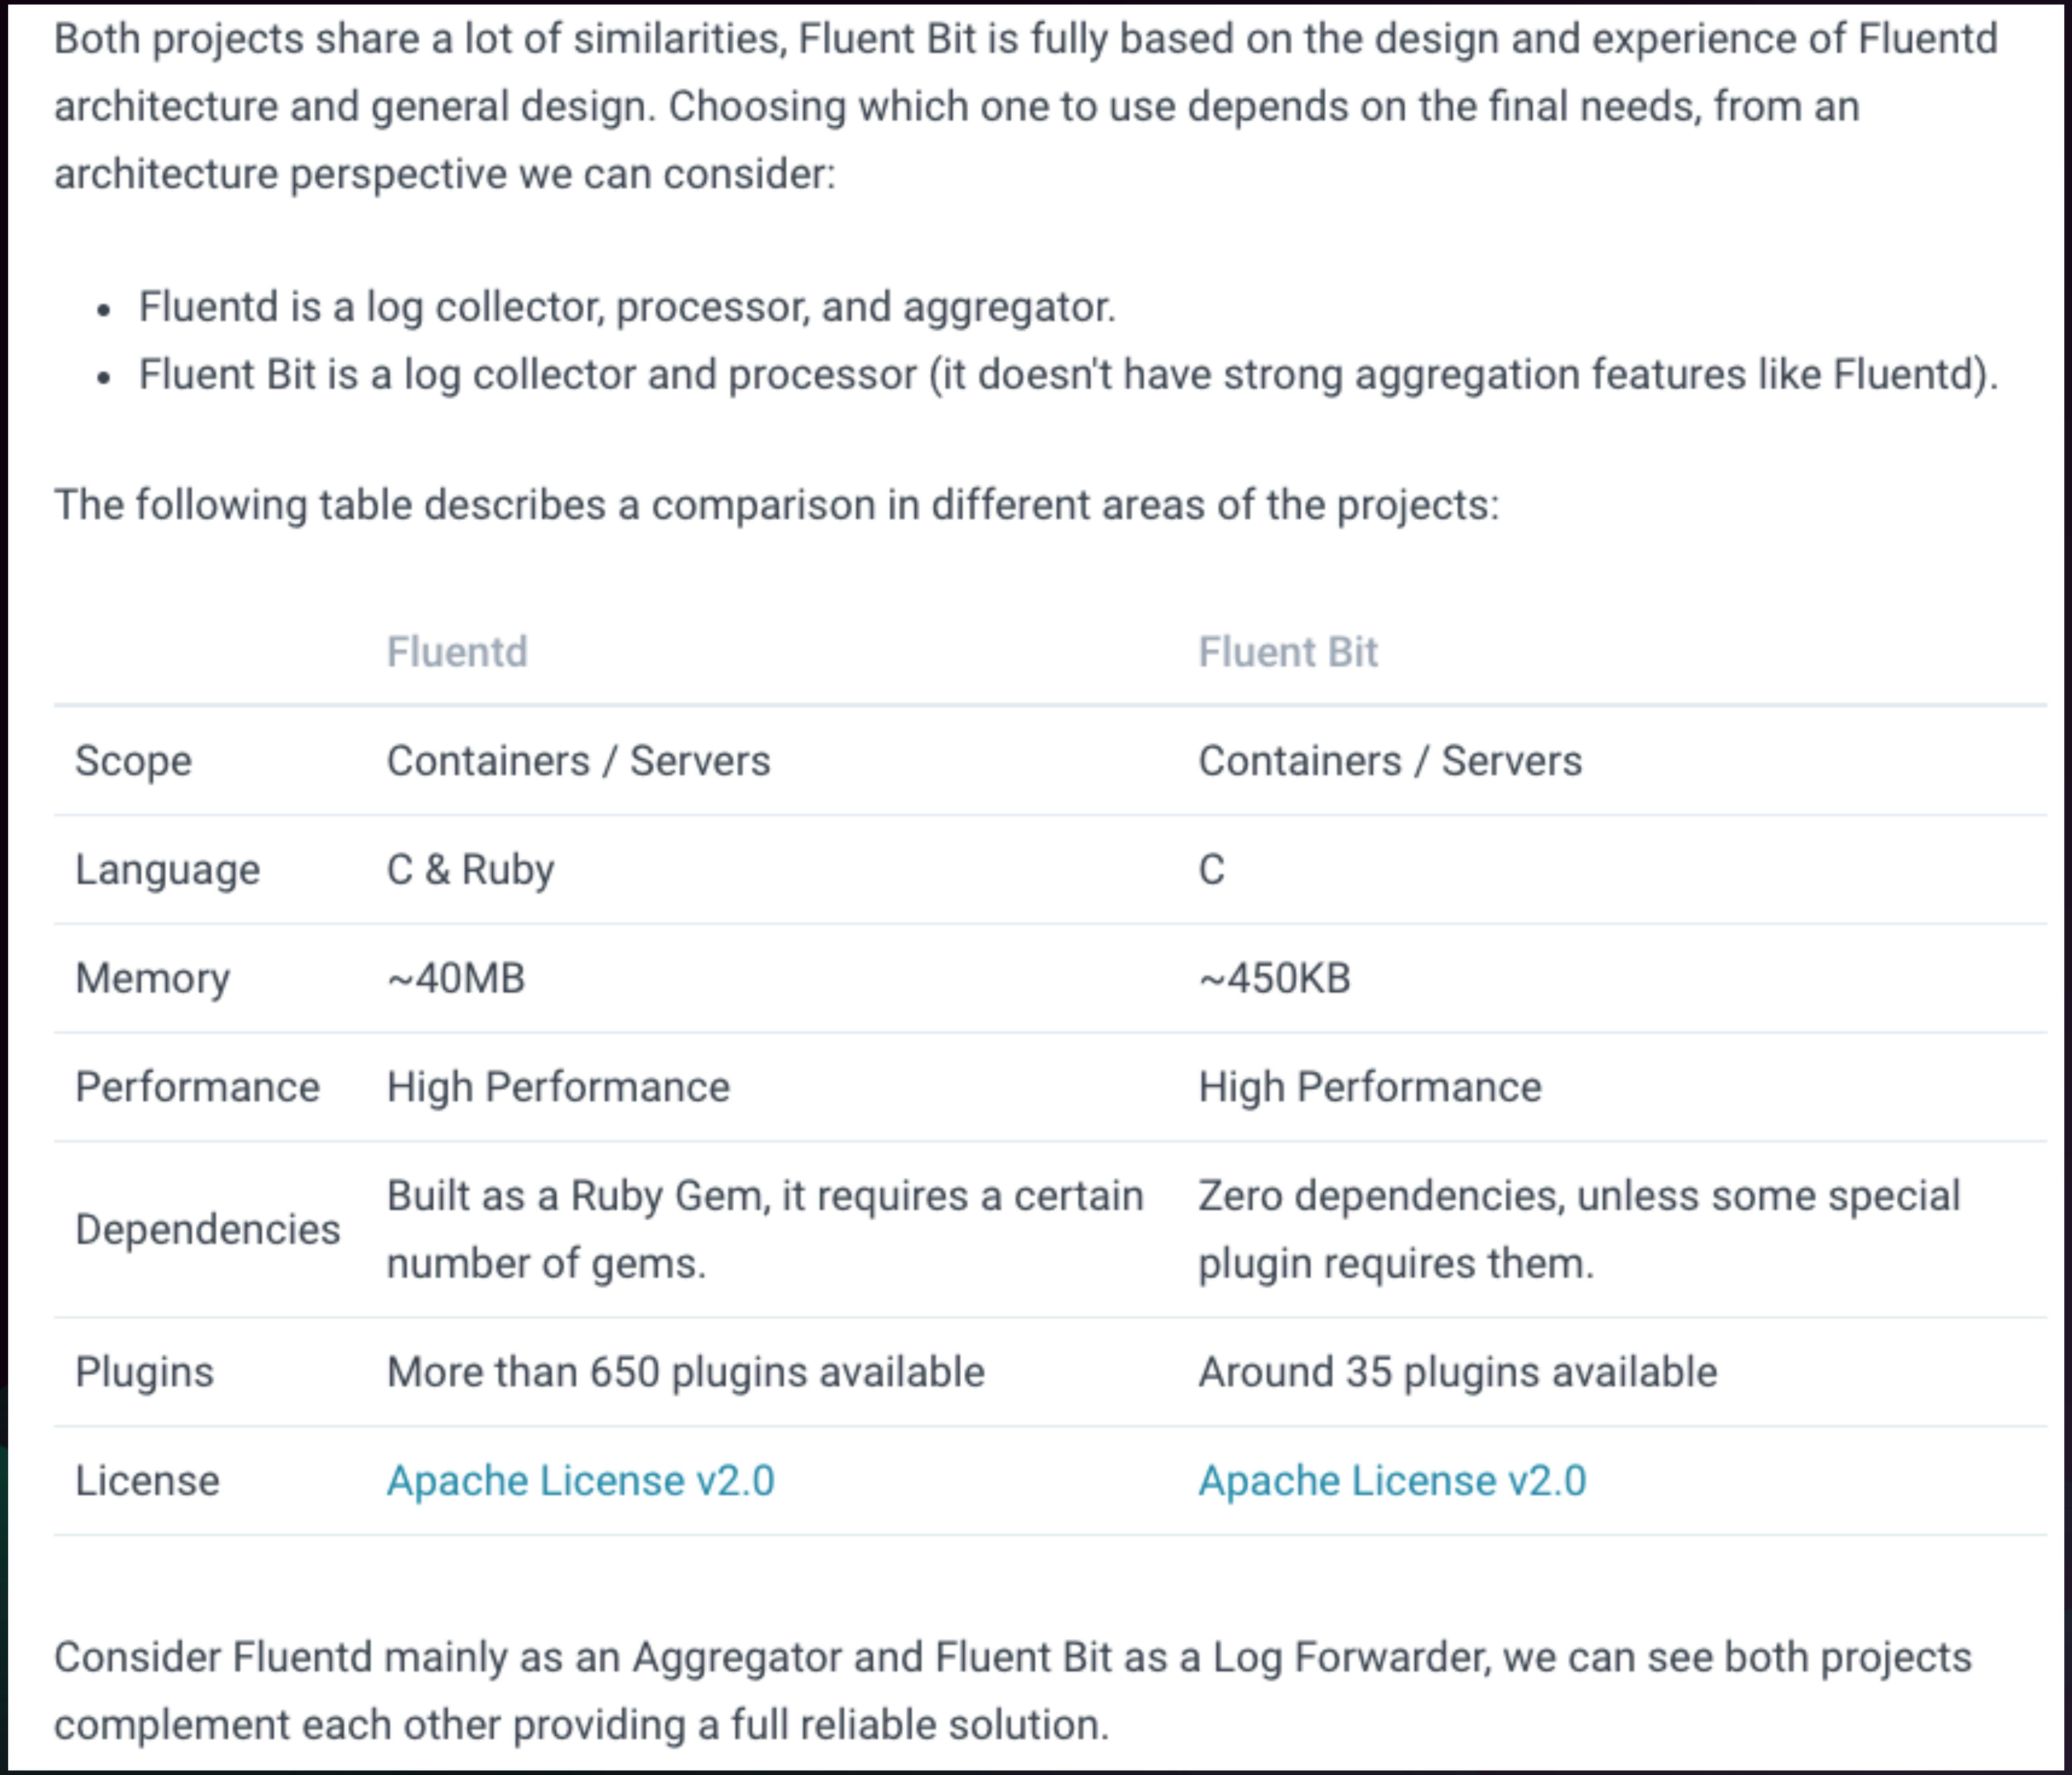
Both projects share a lot of similarities, Fluent Bit is fully based on the design and experience of Fluentd architecture and general design. Choosing which one to use depends on the final needs, from an architecture perspective we can consider:

- Fluentd is a log collector, processor, and aggregator.
- Fluent Bit is a log collector and processor (it doesn't have strong aggregation features like Fluentd).

The following table describes a comparison in different areas of the projects:

	Fluentd	Fluent Bit
Scope	Containers / Servers	Containers / Servers
Language	C & Ruby	C
Memory	~40MB	~450KB
Performance	High Performance	High Performance
Dependencies	Built as a Ruby Gem, it requires a certain number of gems.	Zero dependencies, unless some special plugin requires them.
Plugins	More than 650 plugins available	Around 35 plugins available
License	Apache License v2.0	Apache License v2.0

Consider Fluentd mainly as an Aggregator and Fluent Bit as a Log Forwarder, we can see both projects complement each other providing a full reliable solution.

- https://docs.fluentbit.io/manual/about/fluentd_and_fluentbit
- Resource footprint: 40MB vs 450KB
- Efficiency: C & Ruby vs C
- Dependencies: Ruby gems vs Zero
- Plugins: 650 vs 35

QINGCLOUD | KUBESPHERE

Why Fluent Bit?

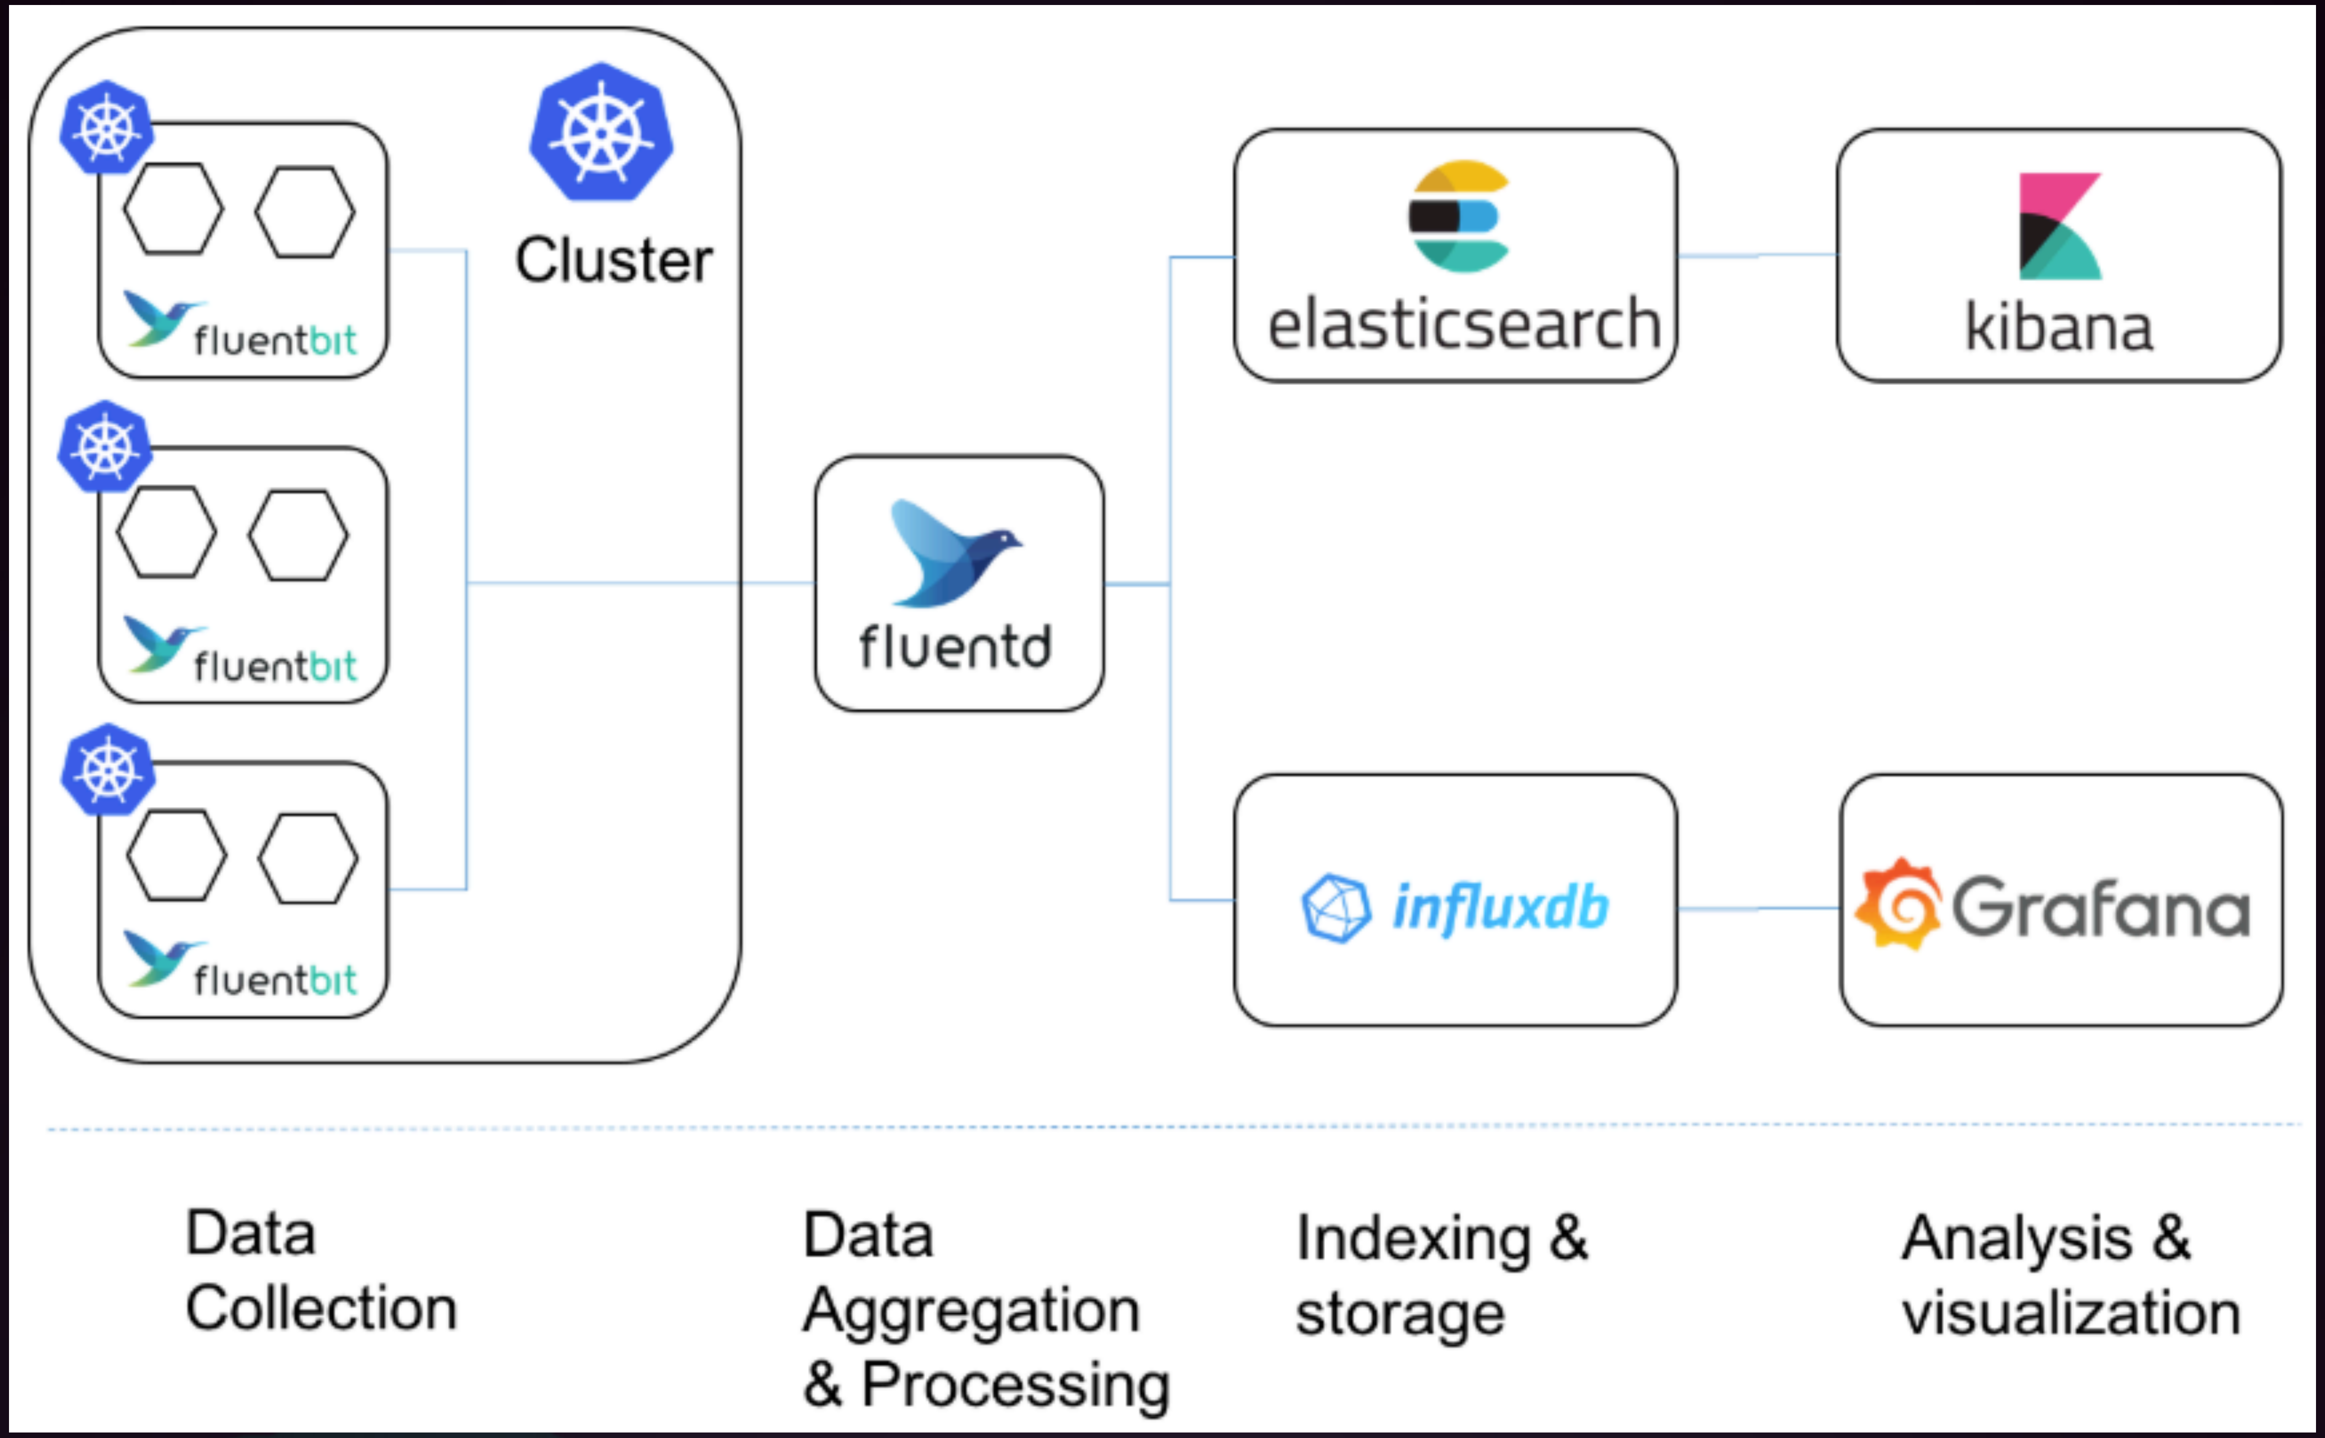

Aggregator vs Collector

Introducing FluentBit Operator

FluentBit can not reload config gracefully

https://github.com/fluent/fluent-bit/pull/842

https://github.com/fluent/fluent-bit/issues/365

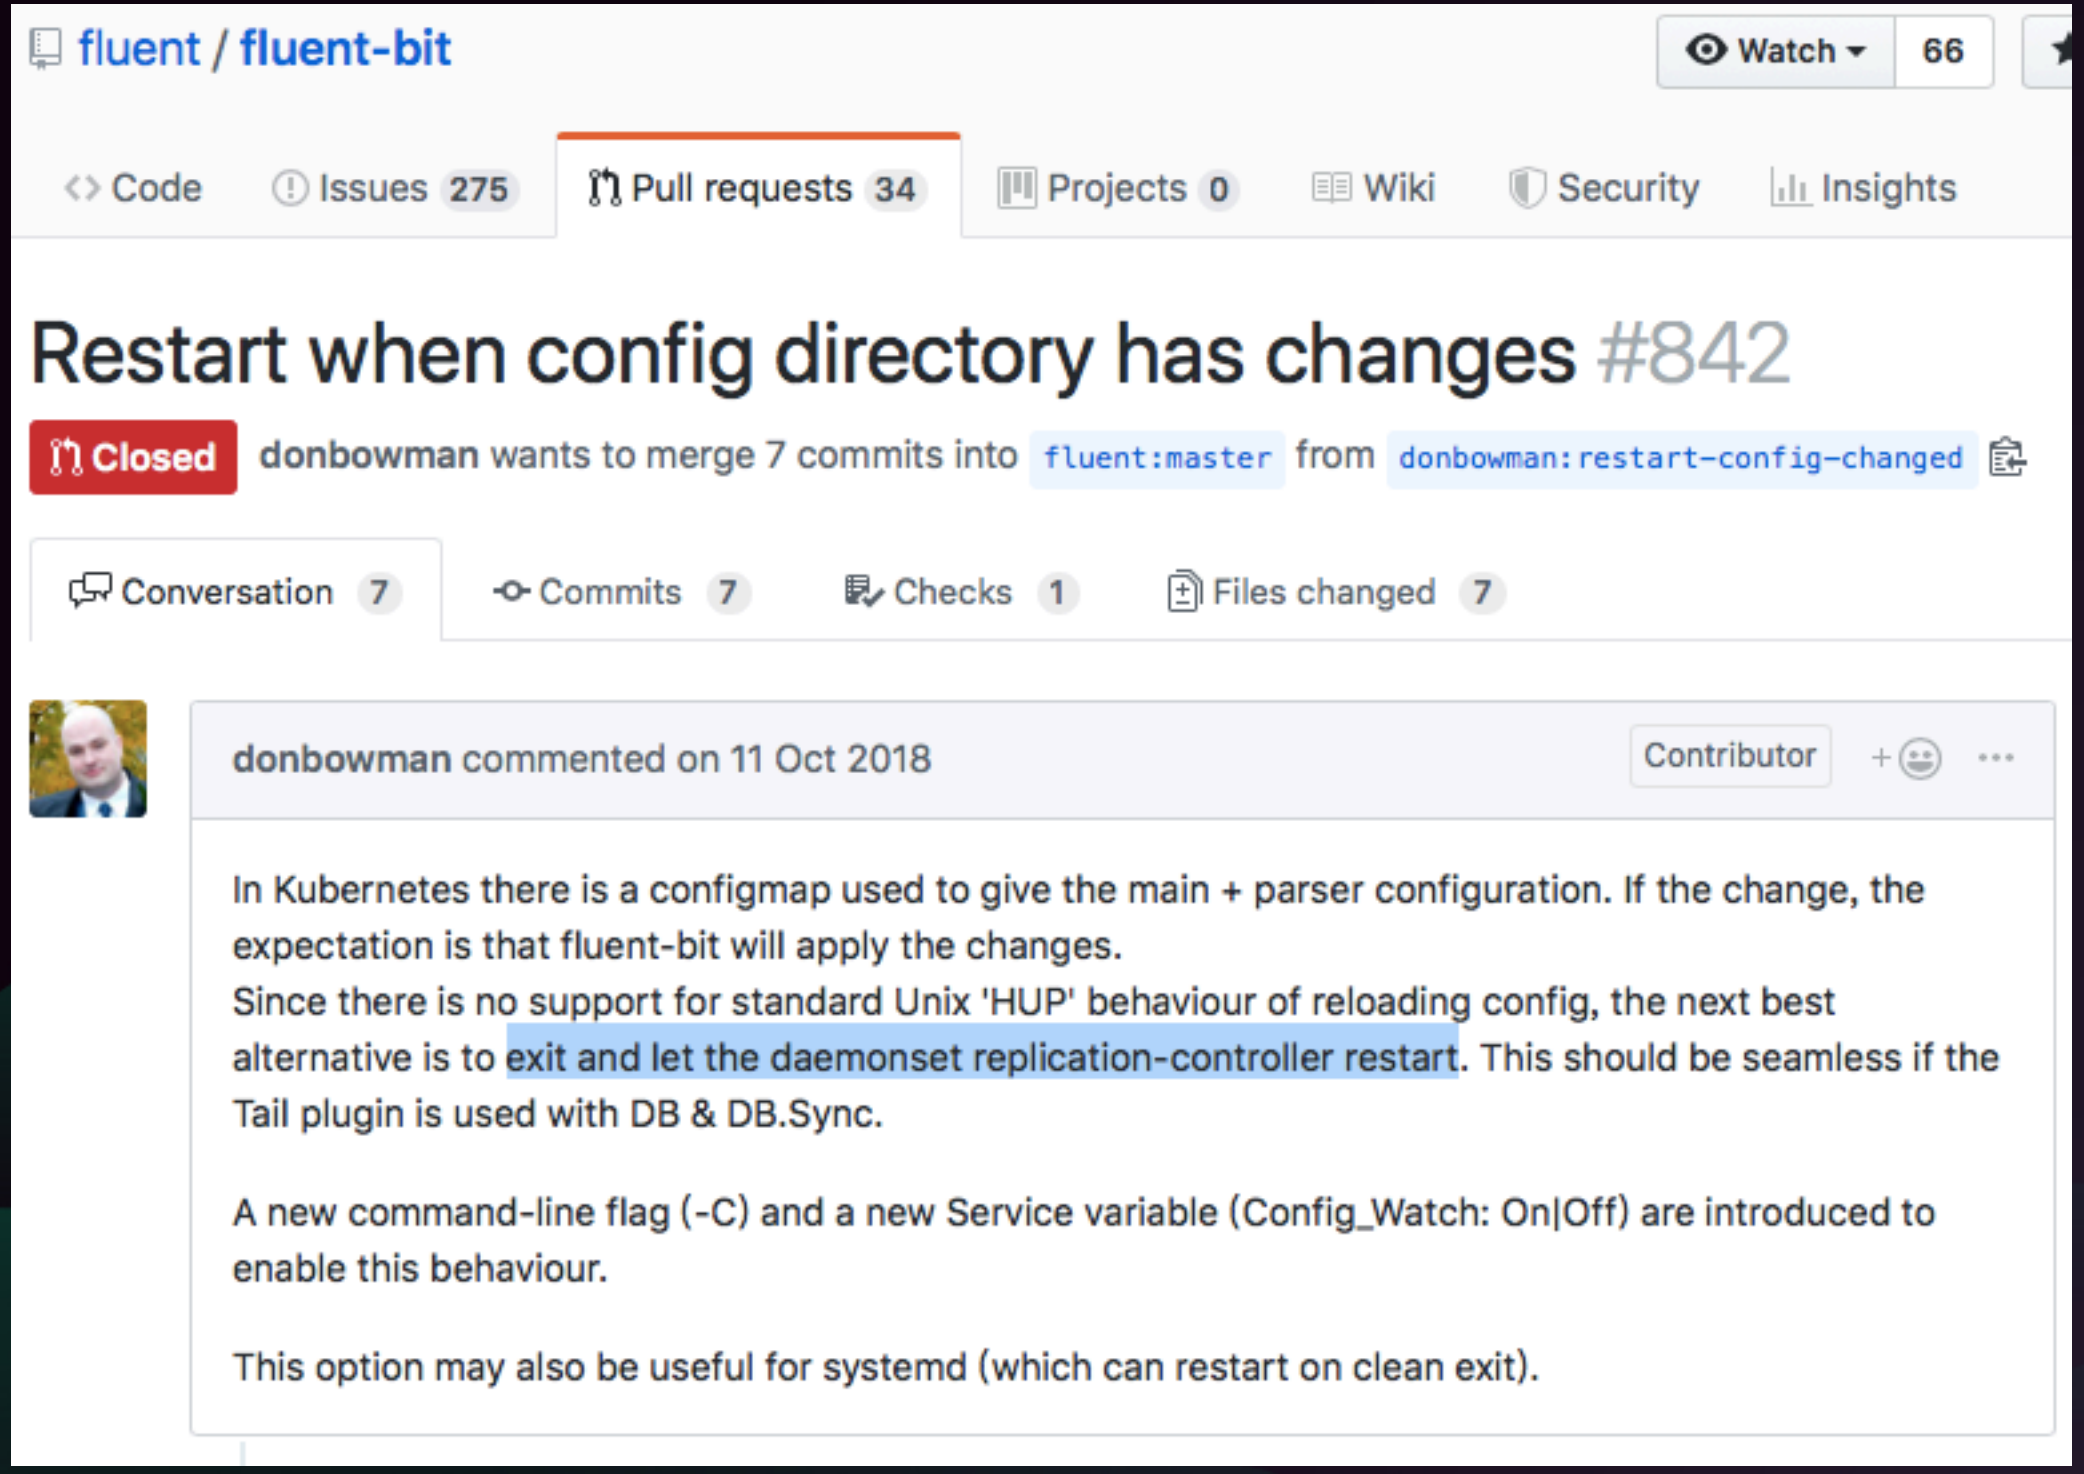

Restart when config directory has changes #842

Closed donbowman wants to merge 7 commits into `fluent:master` from `donbowman:restart-config-changed`

donbowman commented on 11 Oct 2018 Contributor

In Kubernetes there is a configmap used to give the main + parser configuration. If the change, the expectation is that fluent-bit will apply the changes.
Since there is no support for standard Unix 'HUP' behaviour of reloading config, the next best alternative is to exit and let the daemonset replication-controller restart. This should be seamless if the Tail plugin is used with DB & DB.Sync.

A new command-line flag (-C) and a new Service variable (Config_Watch: On|Off) are introduced to enable this behaviour.

This option may also be useful for systemd (which can restart on clean exit).

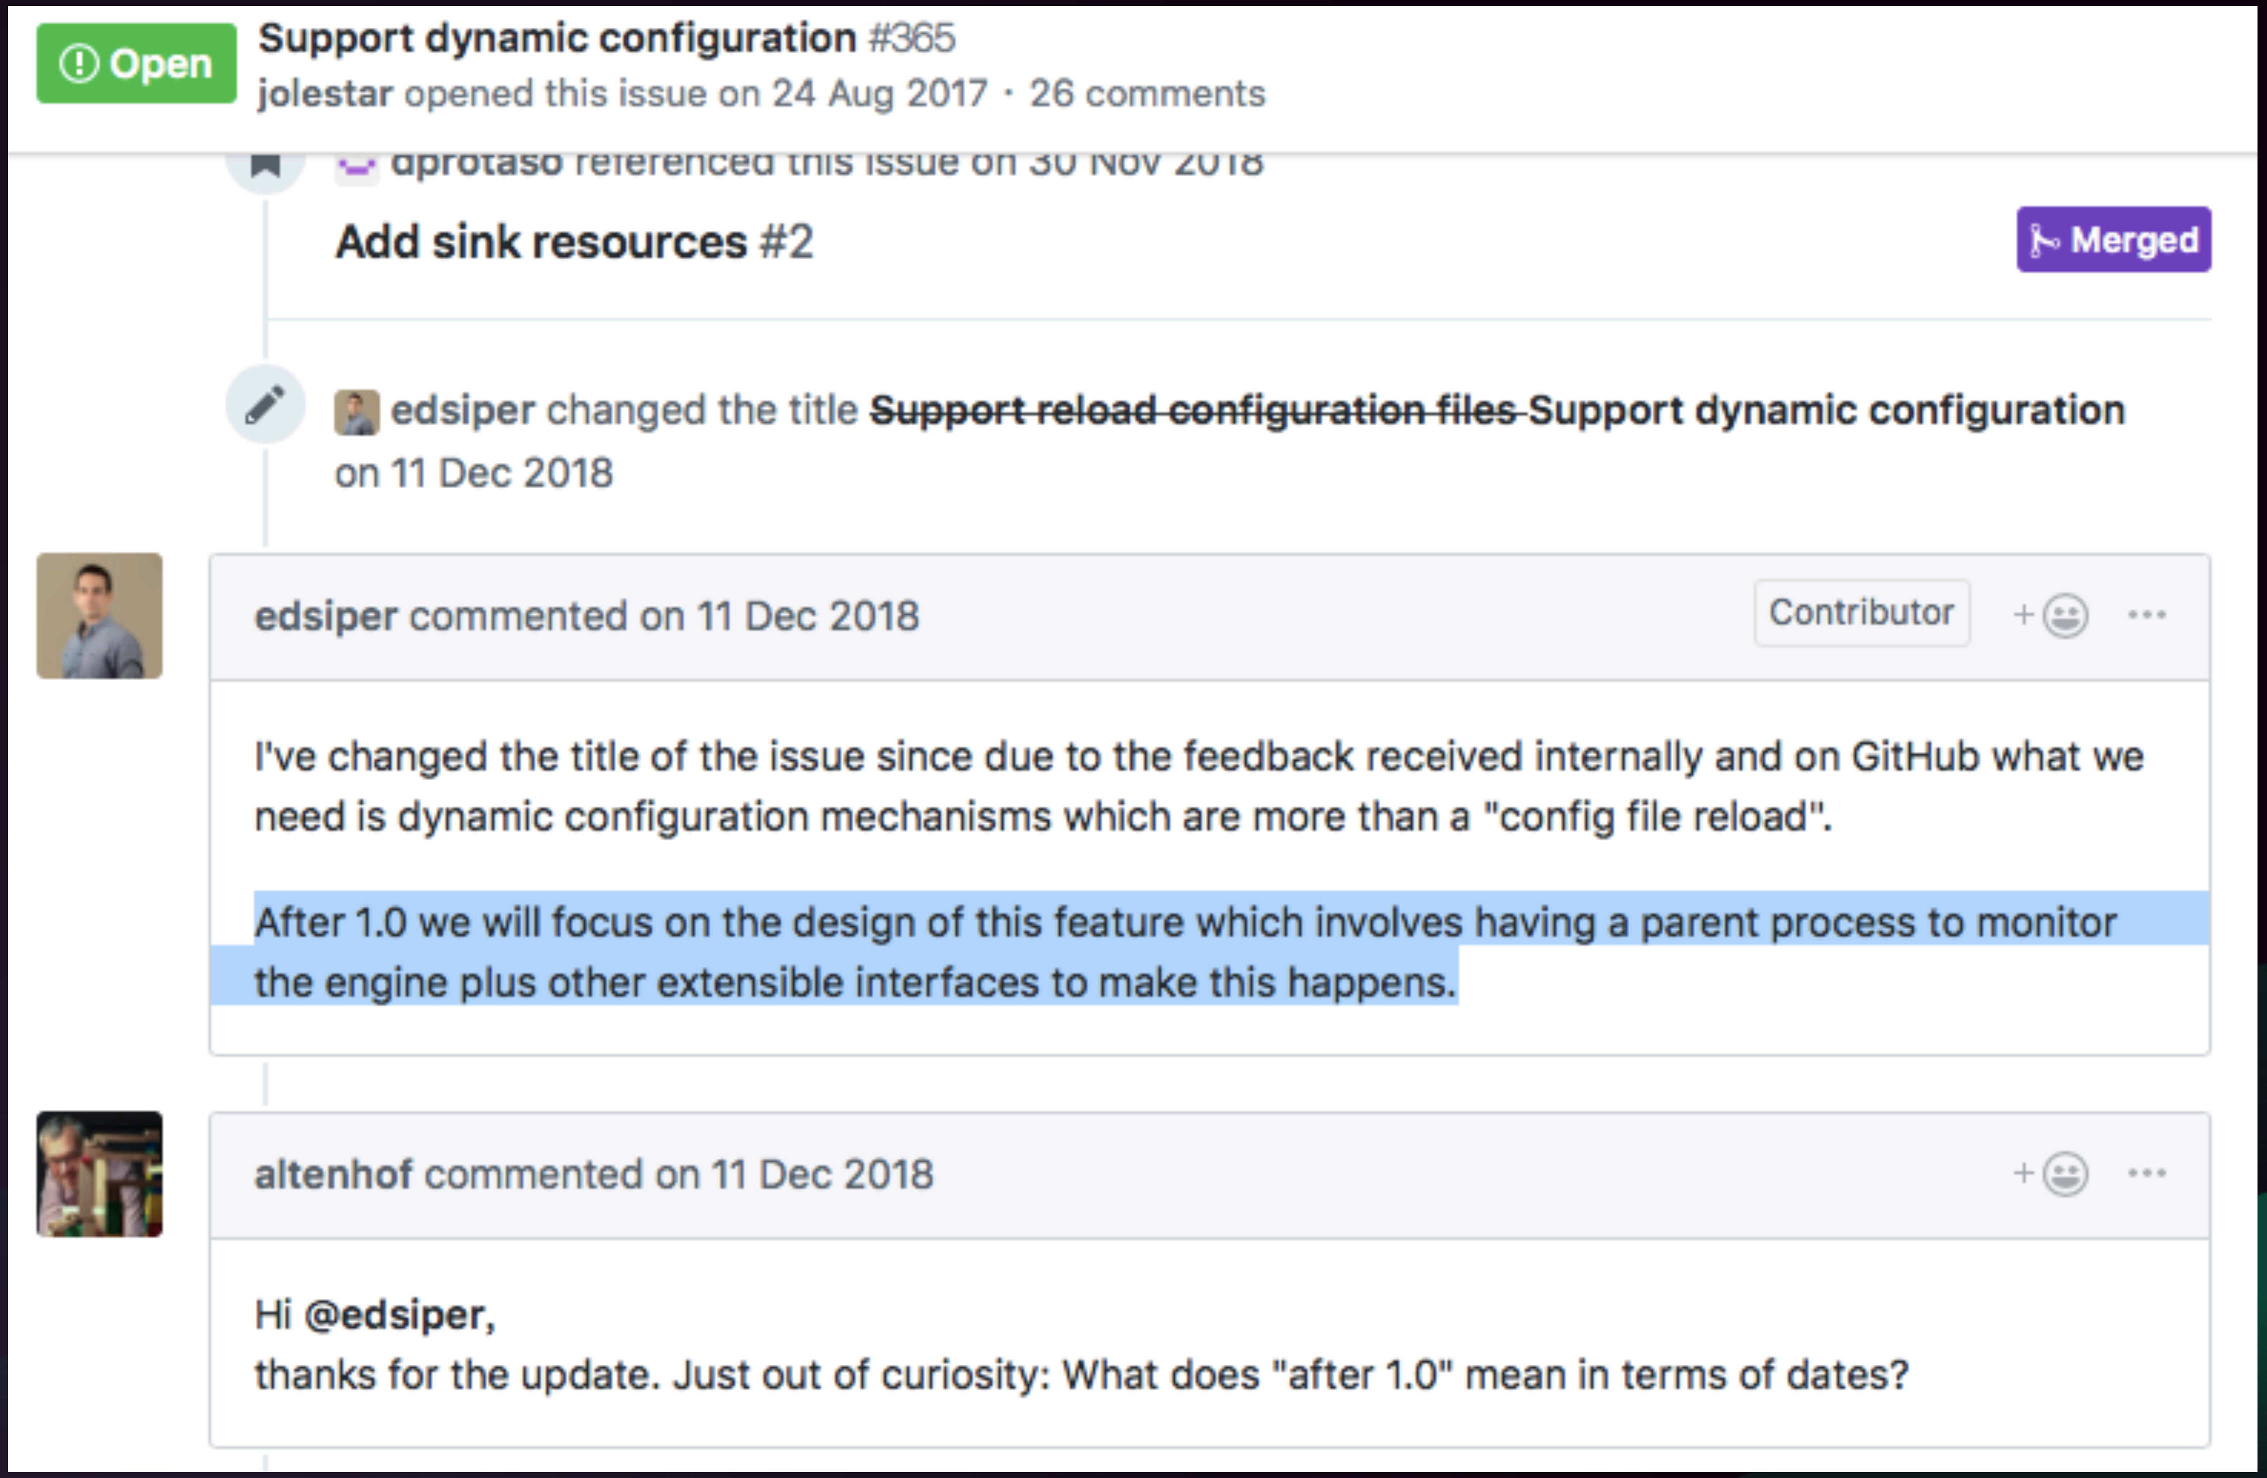

Open Support dynamic configuration #365
jolestar opened this issue on 24 Aug 2017 · 26 comments

aprotaso referenced this issue on 30 Nov 2018

Add sink resources #2 Merged

edsiper changed the title ~~Support reload configuration files~~ Support dynamic configuration on 11 Dec 2018

edsiper commented on 11 Dec 2018 Contributor

I've changed the title of the issue since due to the feedback received internally and on GitHub what we need is dynamic configuration mechanisms which are more than a "config file reload".

After 1.0 we will focus on the design of this feature which involves having a parent process to monitor the engine plus other extensible interfaces to make this happens.

altenhof commented on 11 Dec 2018

Hi @edsiper,
thanks for the update. Just out of curiosity: What does "after 1.0" mean in terms of dates?

QINGCLOUD | KUBESPHERE

Introducing FluentBit Operator

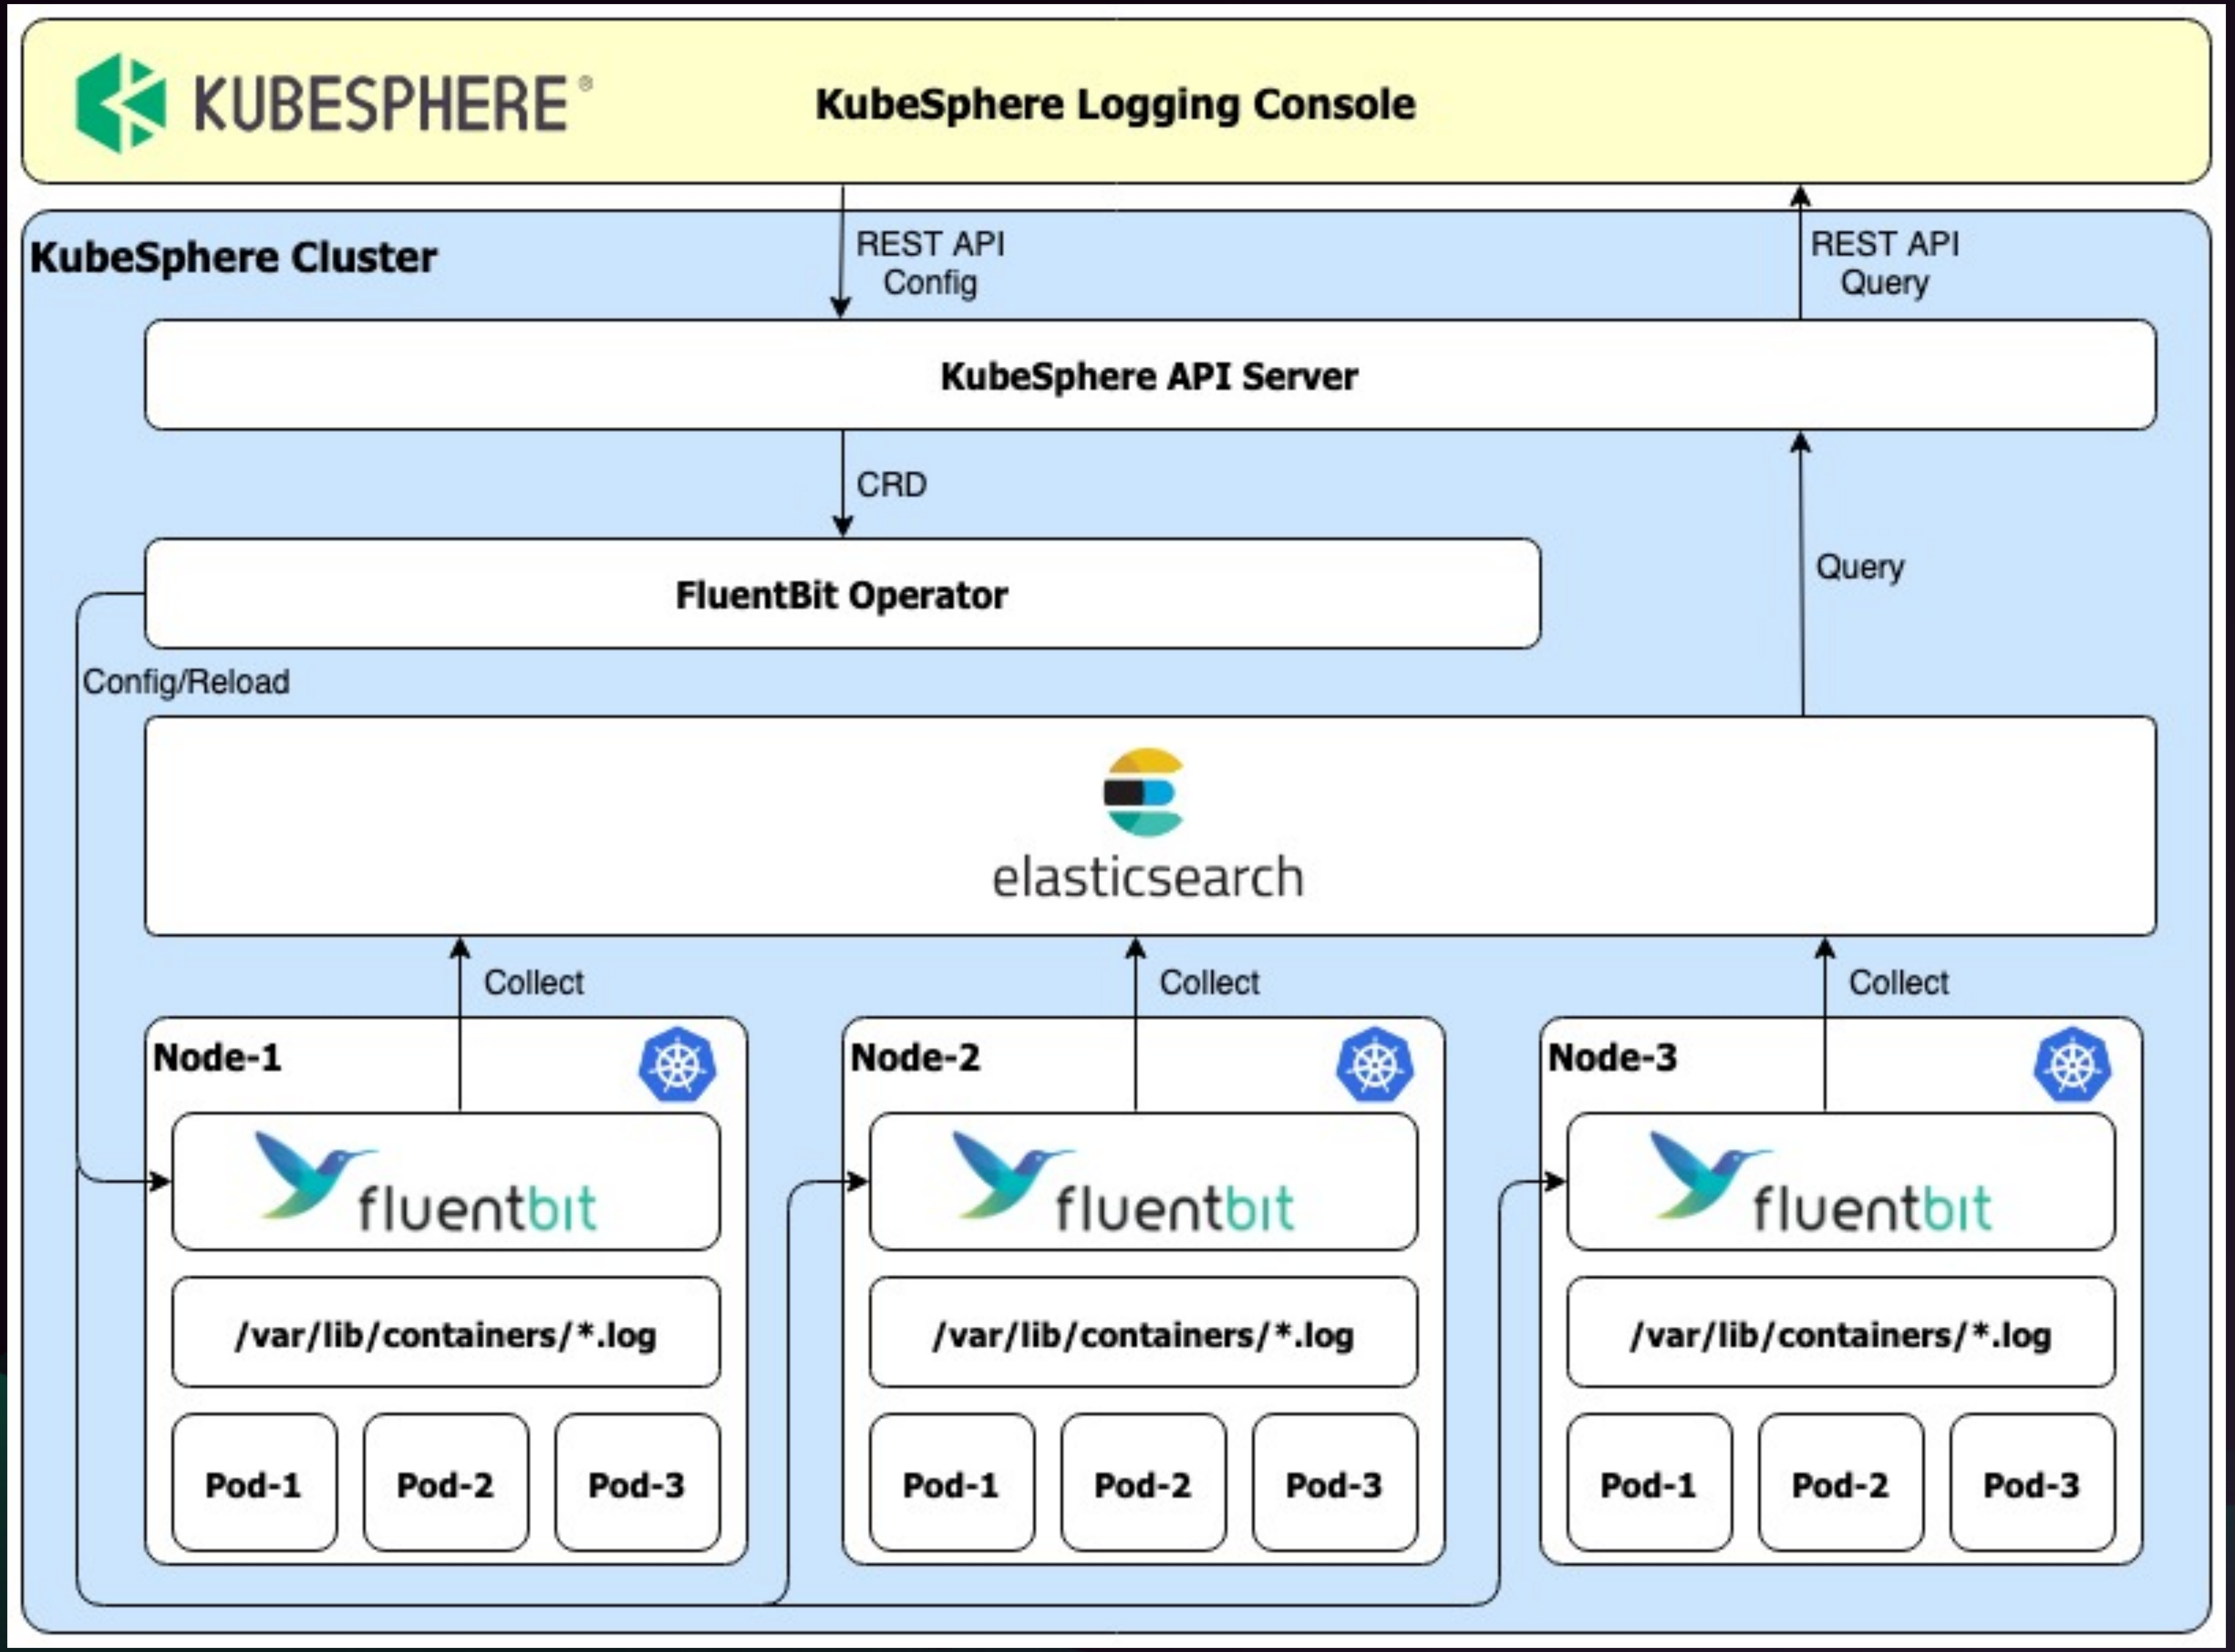

Why FluentBit Operator?

- Config output, filter from UI
- Standard CRD process
  kubectl edit fluentbit fluent-bit
- Reload FluentBit without pod restarting

QINGCLOUD | KUBESPHERE

Introducing FluentBit Operator

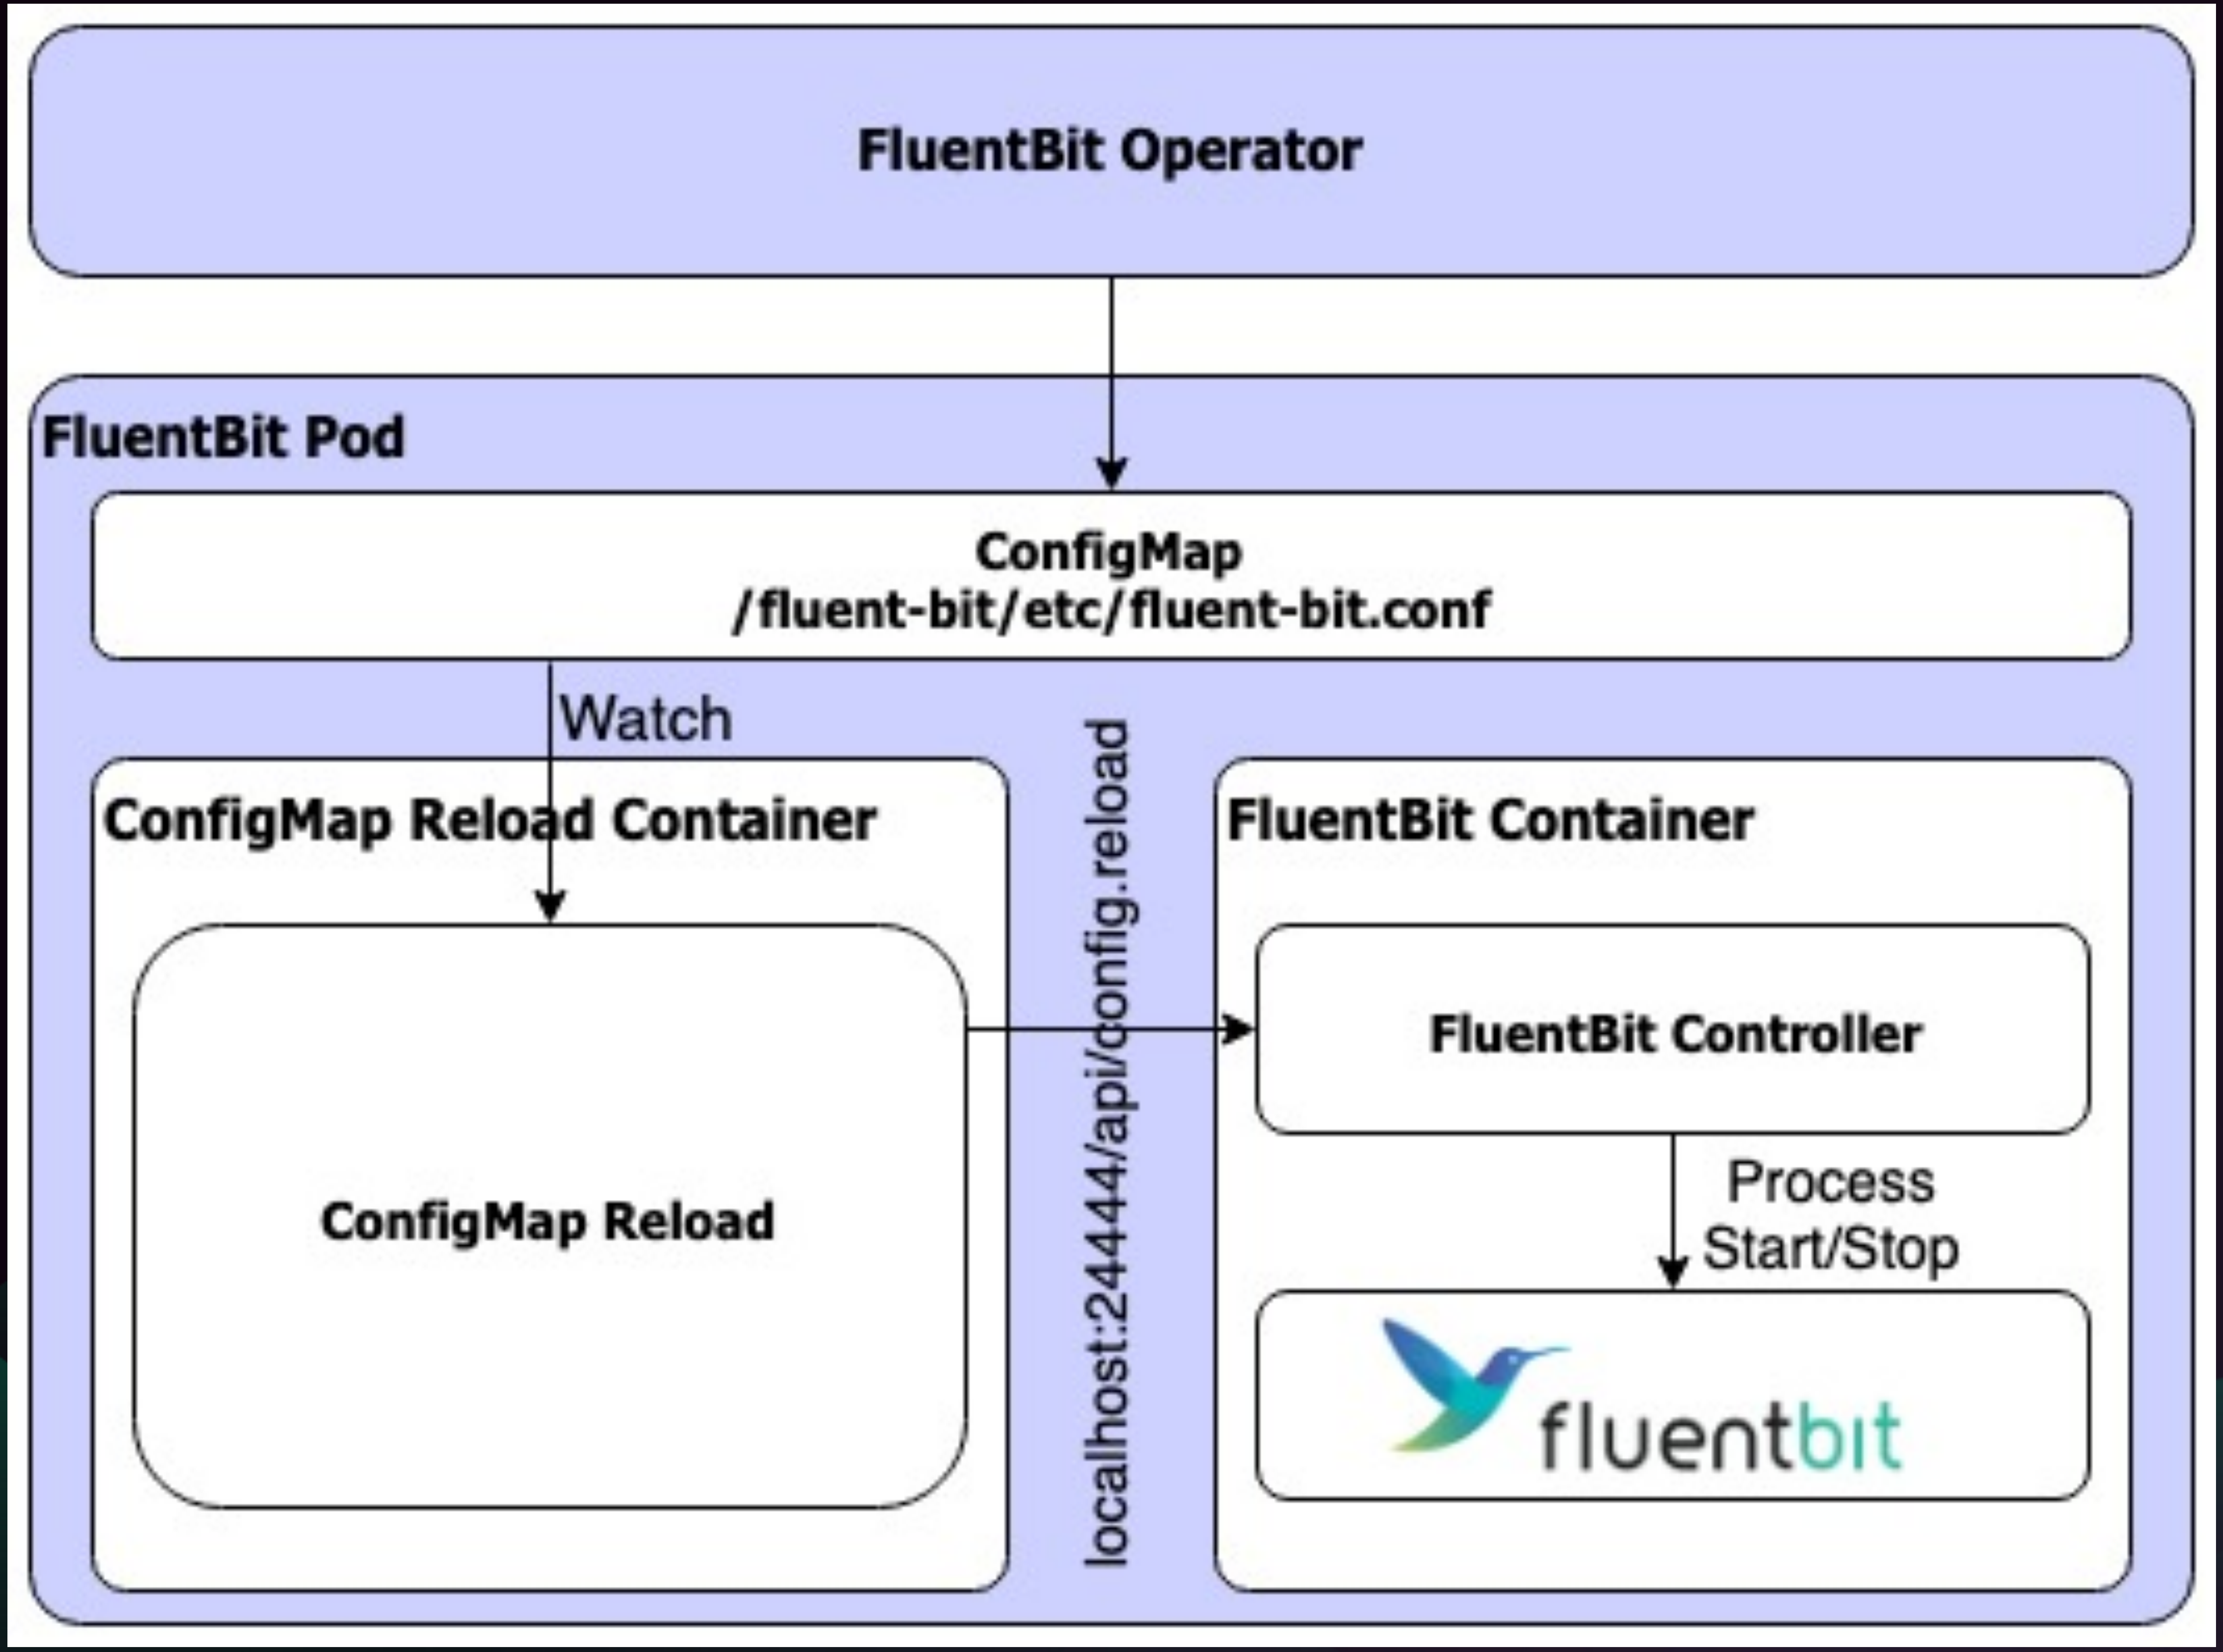

How to reload FluentBit config gracefully?

- Configmap reload sidecar https://github.com/jimmidyson/configmap-reload
- Customized reload interface: http://localhost:24444/api/config.reload

Introducing KubeSphere logging console

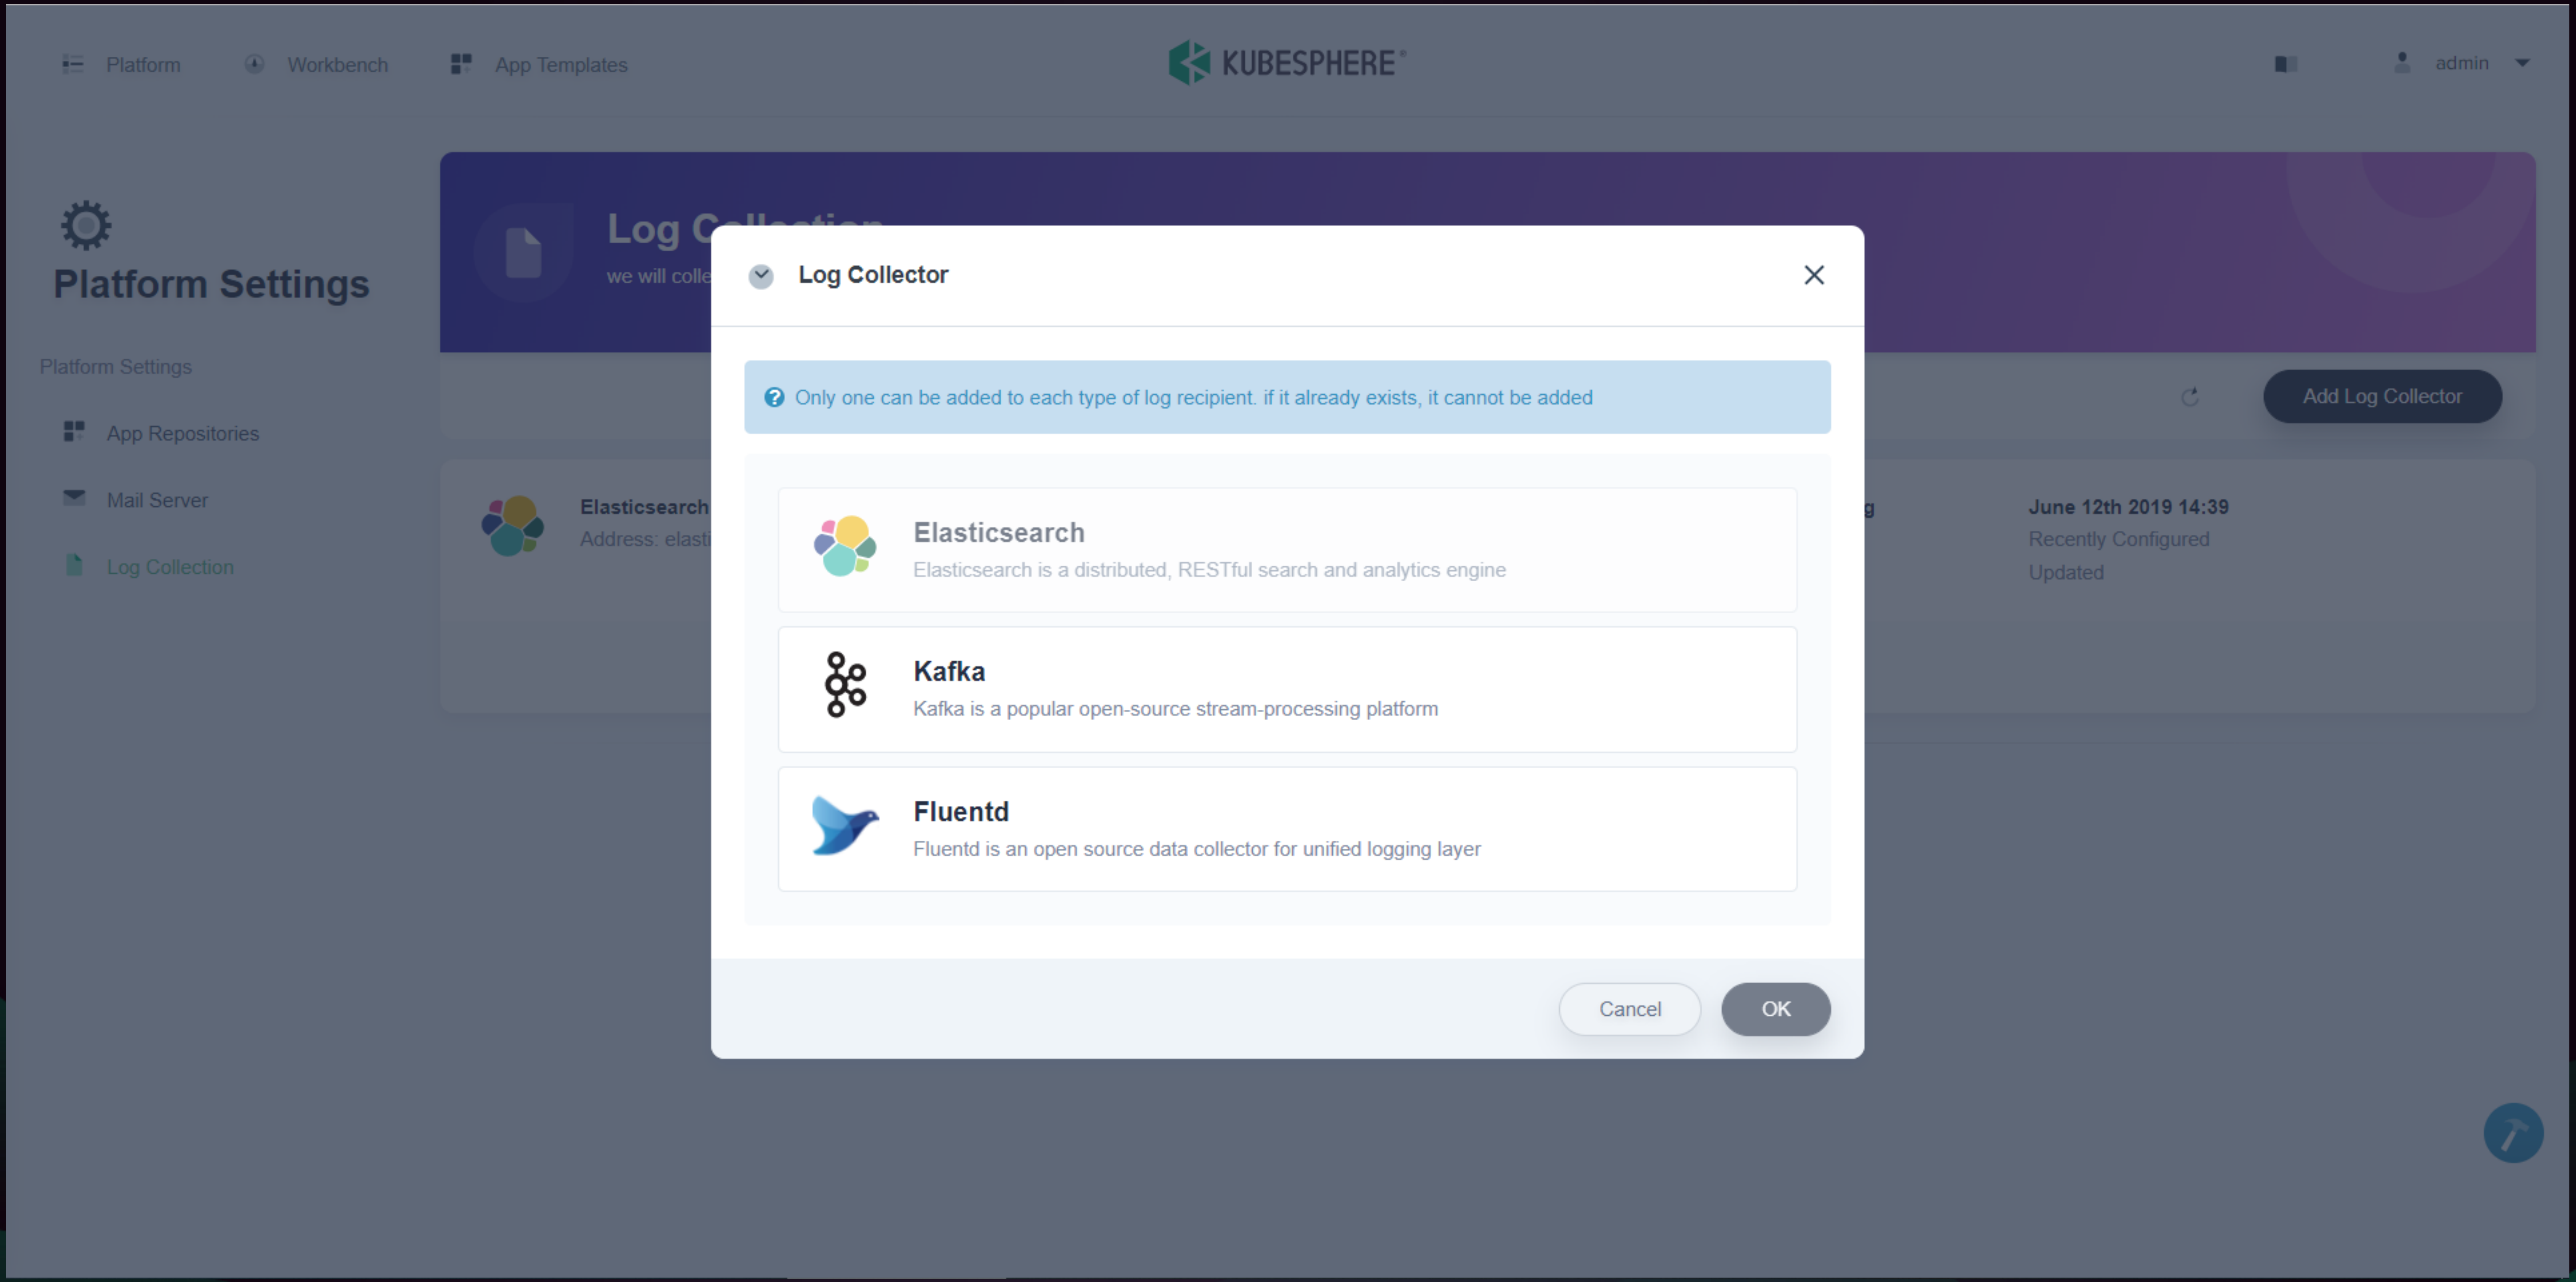

Log Collector

Only one can be added to each type of log recipient. if it already exists, it cannot be added

Elasticsearch
Elasticsearch is a distributed, RESTful search and analytics engine

Kafka
Kafka is a popular open-source stream-processing platform

Fluentd
Fluentd is an open source data collector for unified logging layer

Cancel OK

K8S log config backed by FluentBit Operator

QINGCLOUD | KUBESPHERE

Introducing KubeSphere logging console

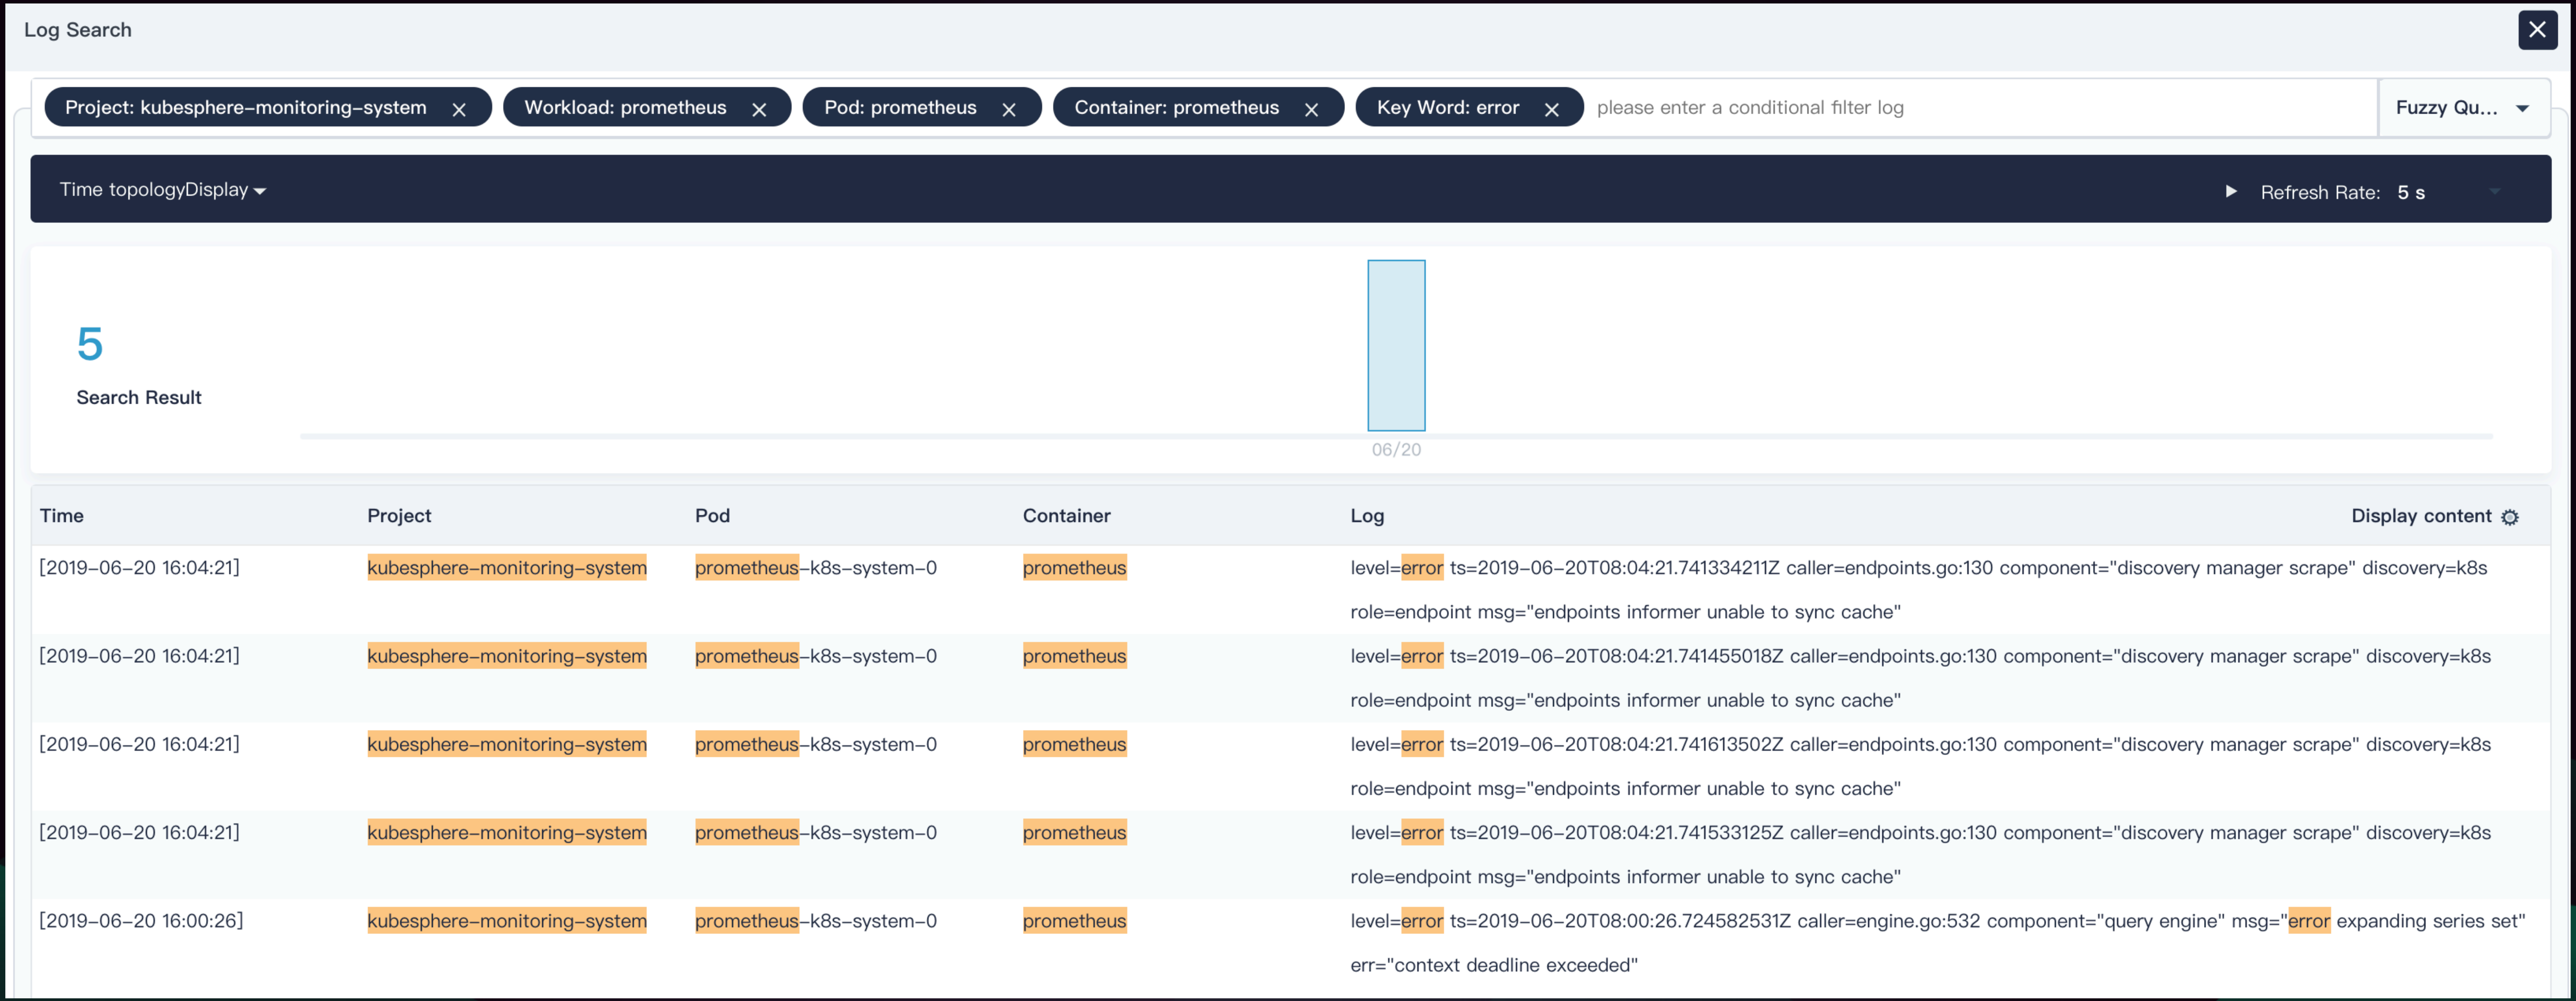

Log Search

Project: kubesphere-monitoring-system ✕ | Workload: prometheus ✕ | Pod: prometheus ✕ | Container: prometheus ✕ | Key Word: error ✕ | please enter a conditional filter log | Fuzzy Qu...

Time topologyDisplay ▾ | Refresh Rate: 5 s

5
Search Result

Time	Project	Pod	Container	Log
[2019-06-20 16:04:21]	kubesphere-monitoring-system	prometheus-k8s-system-0	prometheus	level=error ts=2019-06-20T08:04:21.741334211Z caller=endpoints.go:130 component="discovery manager scrape" discovery=k8s role=endpoint msg="endpoints informer unable to sync cache"
[2019-06-20 16:04:21]	kubesphere-monitoring-system	prometheus-k8s-system-0	prometheus	level=error ts=2019-06-20T08:04:21.741455018Z caller=endpoints.go:130 component="discovery manager scrape" discovery=k8s role=endpoint msg="endpoints informer unable to sync cache"
[2019-06-20 16:04:21]	kubesphere-monitoring-system	prometheus-k8s-system-0	prometheus	level=error ts=2019-06-20T08:04:21.741613502Z caller=endpoints.go:130 component="discovery manager scrape" discovery=k8s role=endpoint msg="endpoints informer unable to sync cache"
[2019-06-20 16:04:21]	kubesphere-monitoring-system	prometheus-k8s-system-0	prometheus	level=error ts=2019-06-20T08:04:21.741533125Z caller=endpoints.go:130 component="discovery manager scrape" discovery=k8s role=endpoint msg="endpoints informer unable to sync cache"
[2019-06-20 16:00:26]	kubesphere-monitoring-system	prometheus-k8s-system-0	prometheus	level=error ts=2019-06-20T08:00:26.724582531Z caller=engine.go:532 component="query engine" msg="error expanding series set" err="context deadline exceeded"

Display content

K8S native log search

QINGCLOUD | KUBESPHERE

Introducing KubeSphere logging console

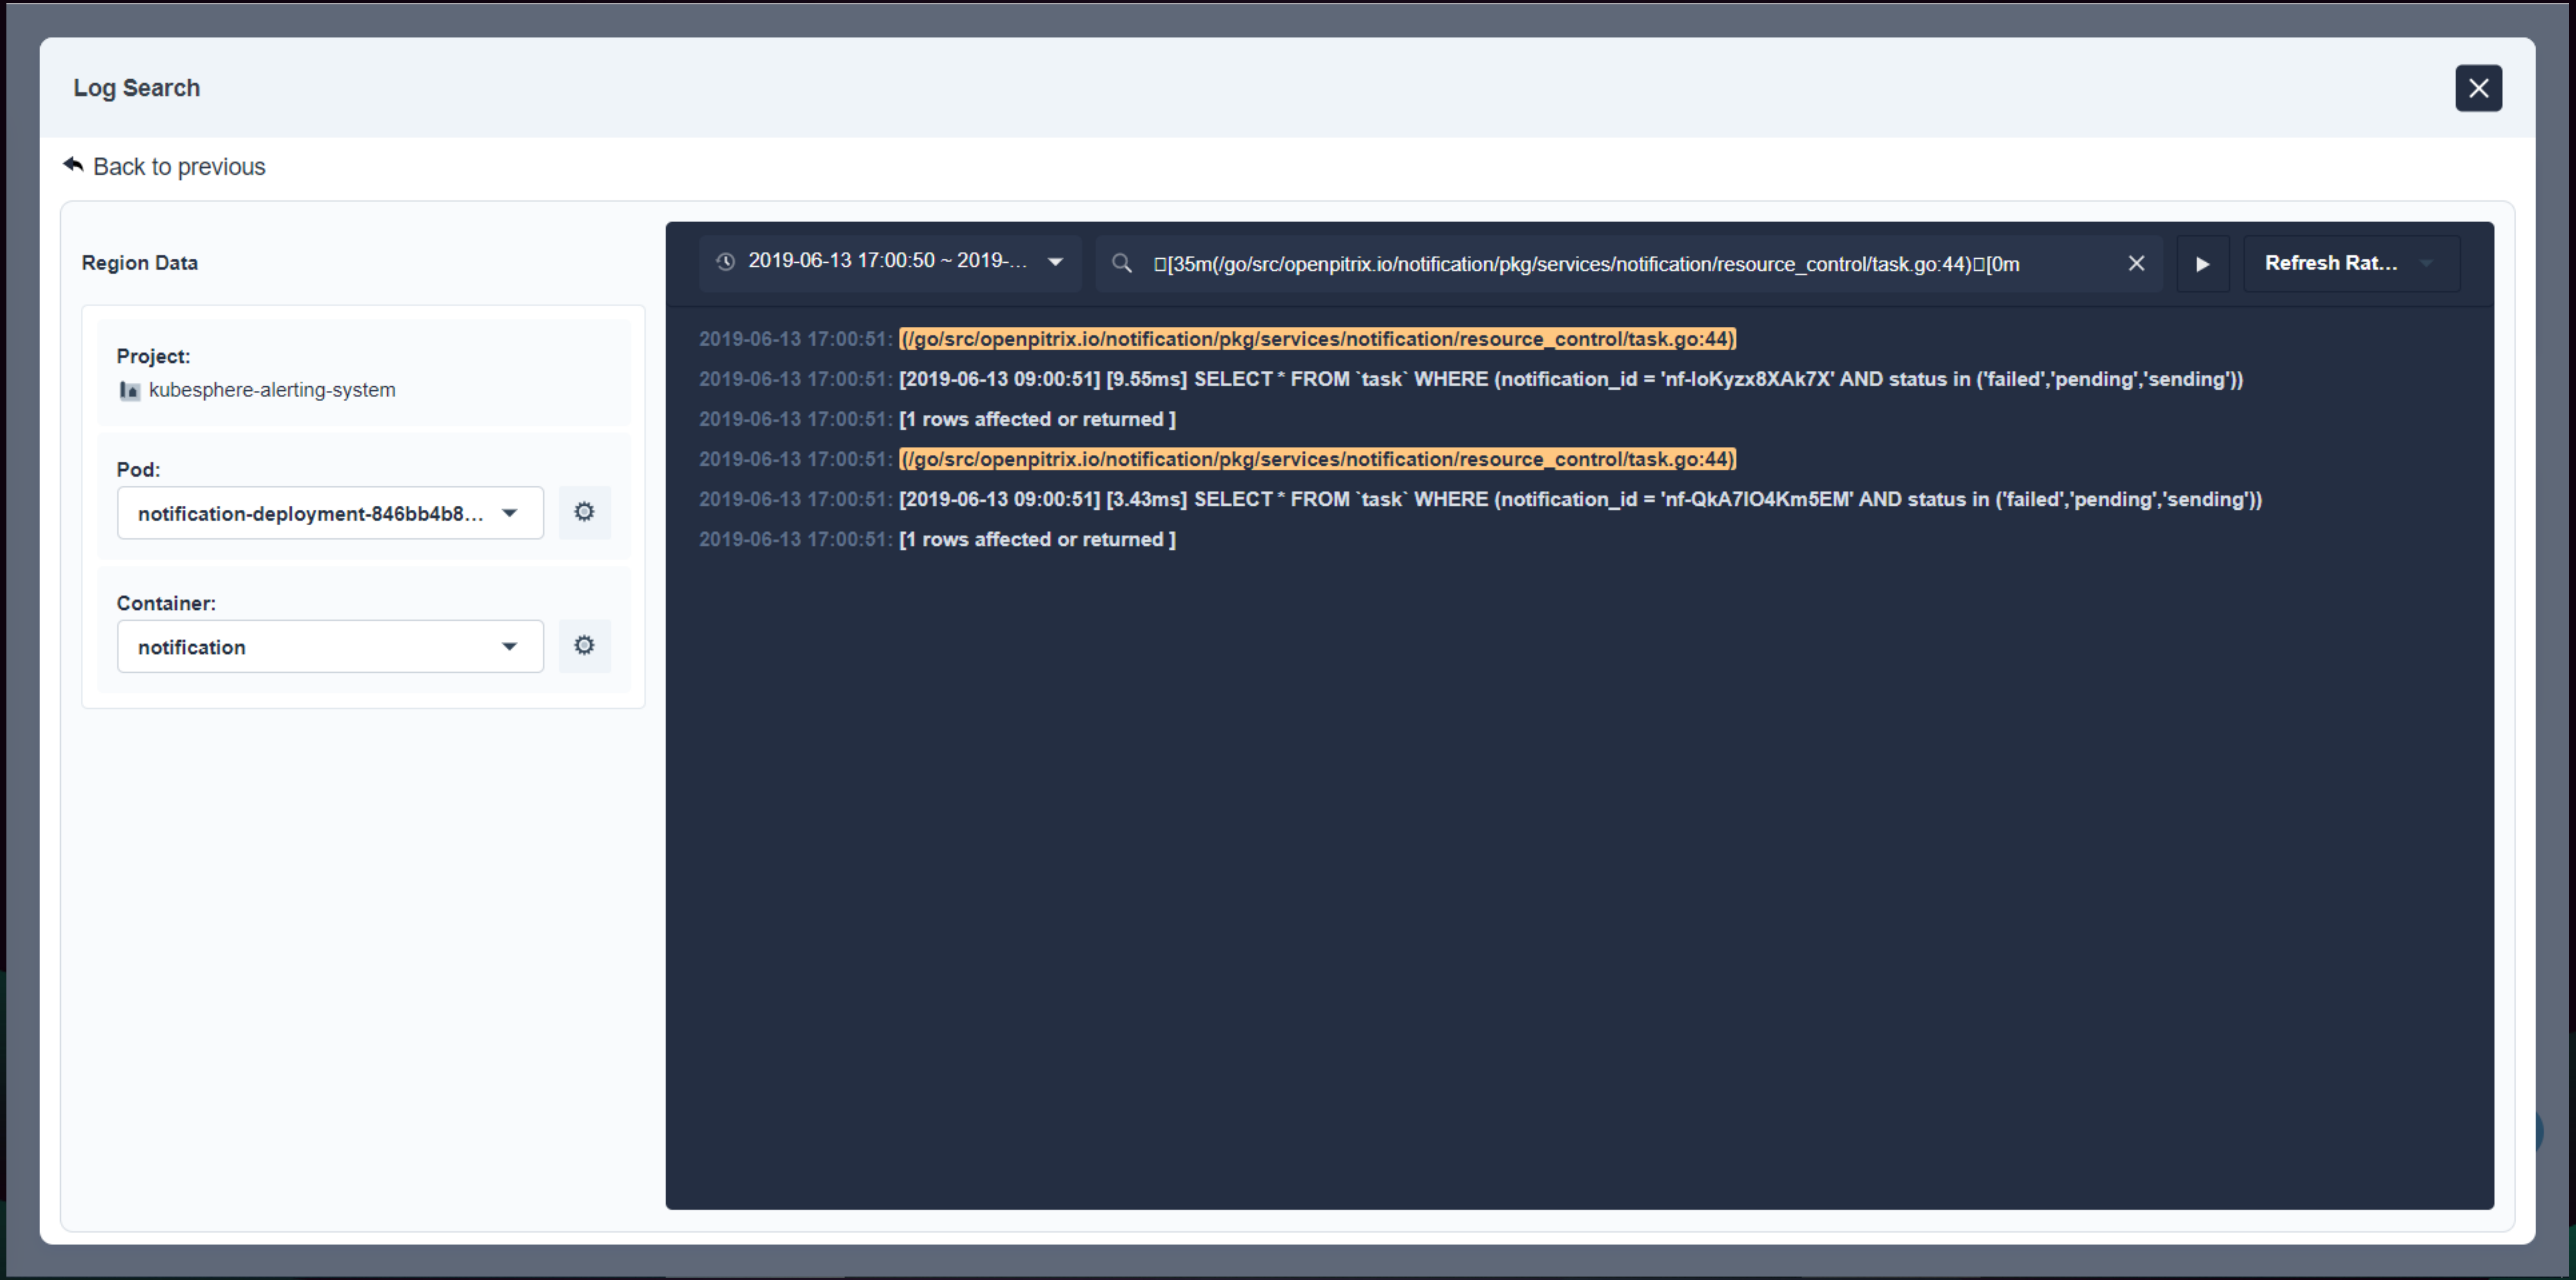

Project: kubesphere-alerting-system

Pod: notification-deployment-846bb4b8...

Container: notification

2019-06-13 17:00:50 ~ 2019-...

□[35m(/go/src/openpitrix.io/notification/pkg/services/notification/resource_control/task.go:44)□[0m

```
2019-06-13 17:00:51: (/go/src/openpitrix.io/notification/pkg/services/notification/resource_control/task.go:44)
2019-06-13 17:00:51: [2019-06-13 09:00:51] [9.55ms] SELECT * FROM `task` WHERE (notification_id = 'nf-loKyzx8XAk7X' AND status in ('failed','pending','sending'))
2019-06-13 17:00:51: [1 rows affected or returned ]
2019-06-13 17:00:51: (/go/src/openpitrix.io/notification/pkg/services/notification/resource_control/task.go:44)
2019-06-13 17:00:51: [2019-06-13 09:00:51] [3.43ms] SELECT * FROM `task` WHERE (notification_id = 'nf-QkA7lO4Km5EM' AND status in ('failed','pending','sending'))
2019-06-13 17:00:51: [1 rows affected or returned ]
```

K8S native log details

QINGCLOUD | KUBESPHERE

Demo

We’re open sourcing

- https://kubesphere.io/
- https://github.com/kubesphere/kubesphere
- https://github.com/kubesphere/fluentbit-operator