A Deep Dive Into Kubernetes Schema Validation

datree.io

$ whoami

```
apiVersion: 32
kind: CPO
metadata:
     name: “Eyar Zilberman”
     labels:
          company: Datree
          role: co-founder
more:
  - Organizer of the biggest github community
  - Hate SQL
  - Love RegEx
```

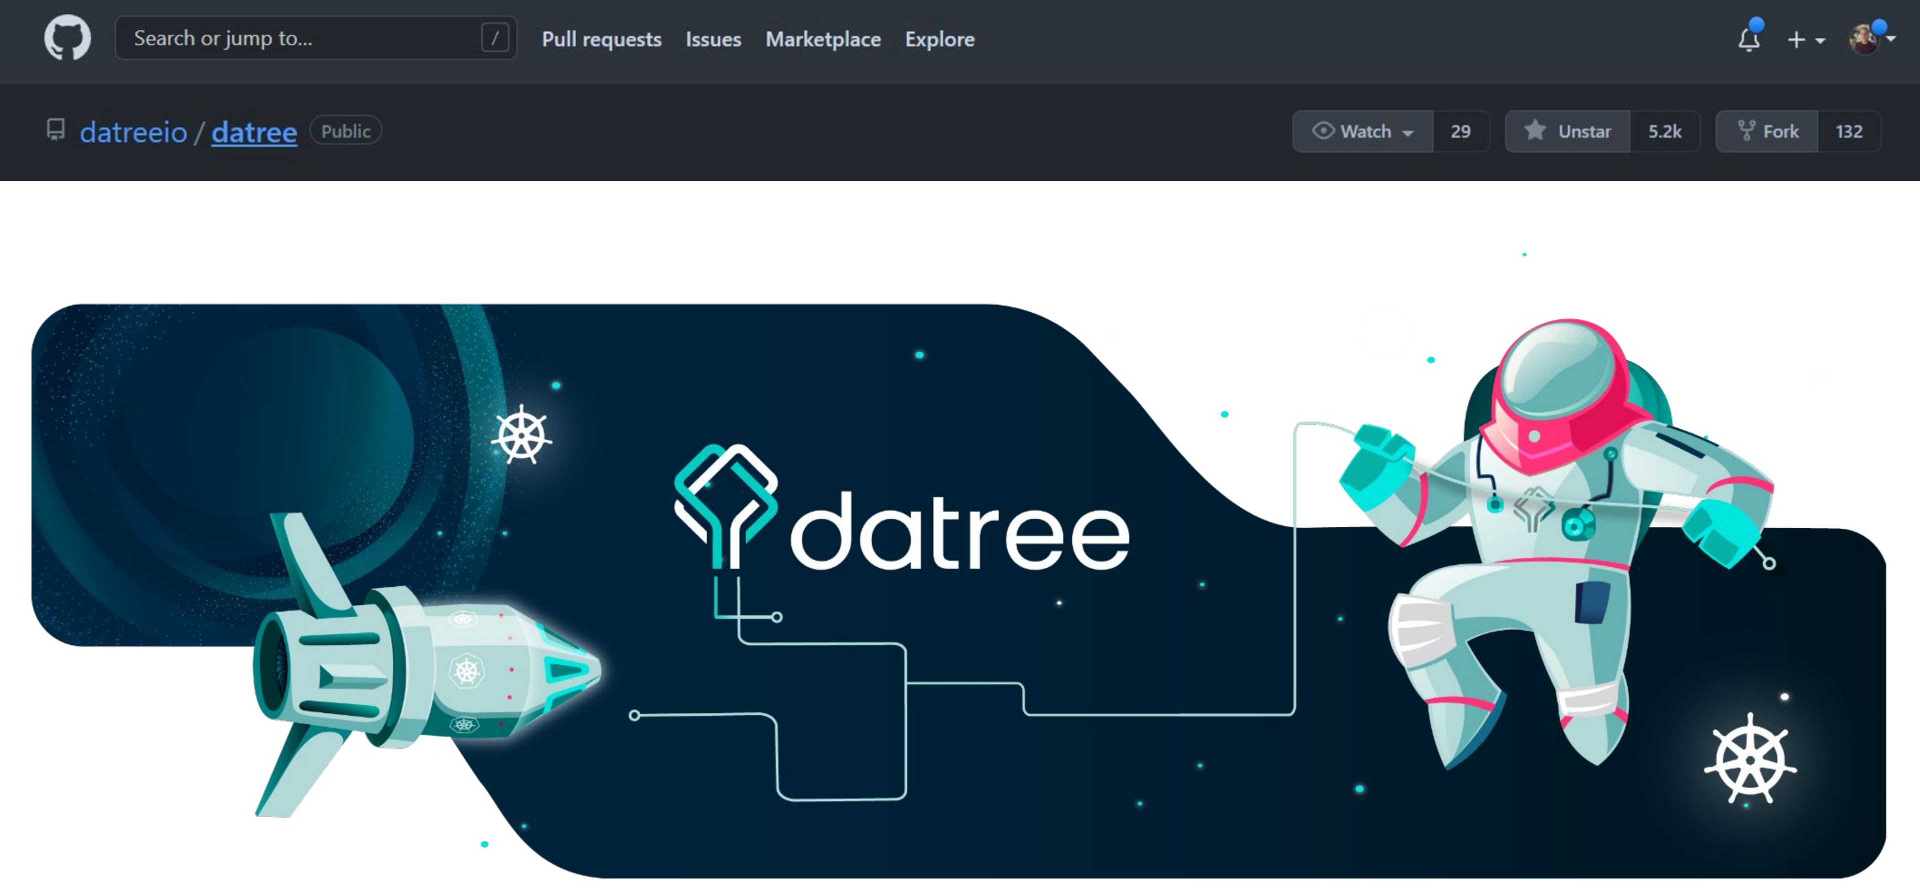

datreeio / datree Public

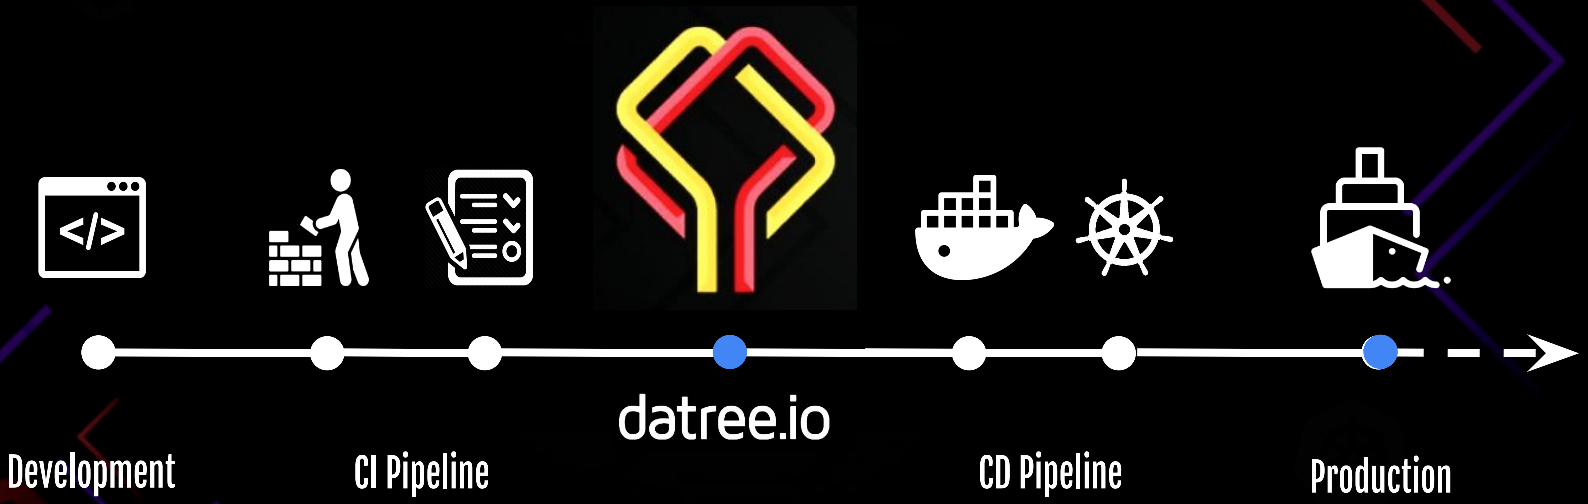

Development

CI Pipeline

datree.io

CD Pipeline

Production

What is Kubernetes Schema validation?

🧪 Set of “unit tests” to verify manifest contains the correct properties (key:value)

[X] K8s.yaml

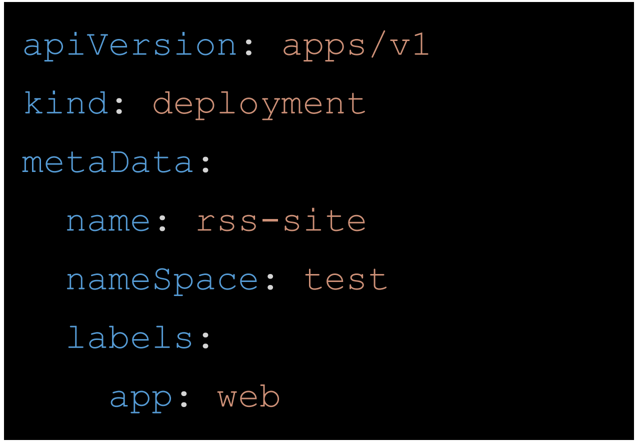

```yaml
apiVersion: apps/v1
kind: deployment
metaData:
  name: rss-site
  nameSpace: test
  labels:
    app: web
```

[V] K8s.yaml

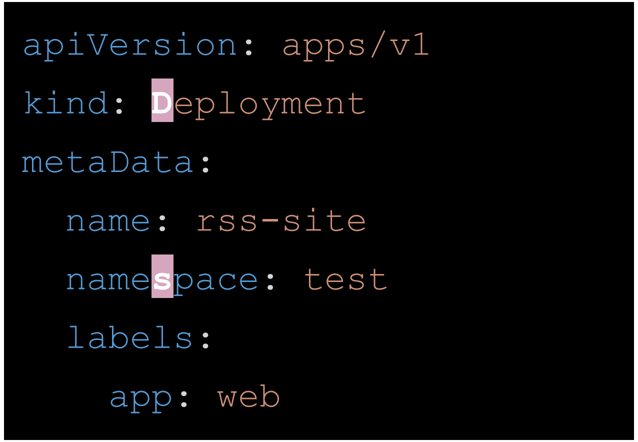

```yaml
apiVersion: apps/v1
kind: Deployment
metaData:
  name: rss-site
  namespace: test
  labels:
    app: web
```

🚢 The schema definition is provided by the community

What is not part of the schema validation?

- YAML syntax validation
  - Correct indentation
- Best practices:
  - Stability - Each container has a configured Memory and CPU limit set
- Team/org policies:
  - Security - Pull all images from private registry (`artifactory.io/nginx:1.16.8`)

How do I use it?

Good news :)

- Activated by default when you apply configs to your cluster

Bad news :(

- It’s too late!

How can check it earlier (“shift-left”)?

- kubectl --dry-run=client/server

kubectl server vs. client

Parameter	Server mode	~~Client mode~~
Preform schema validation?	Yes	~~Yes~~
Preform extra validations?	Yes	
Supports different K8s schema versions?	No	
Requires a connection to your cluster?	Yes	~~Yes~~

BUG: kubernetes/kubernetes/issues/51475

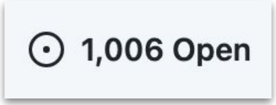

⊙ 1,006 Open

Open-source to the rescue!

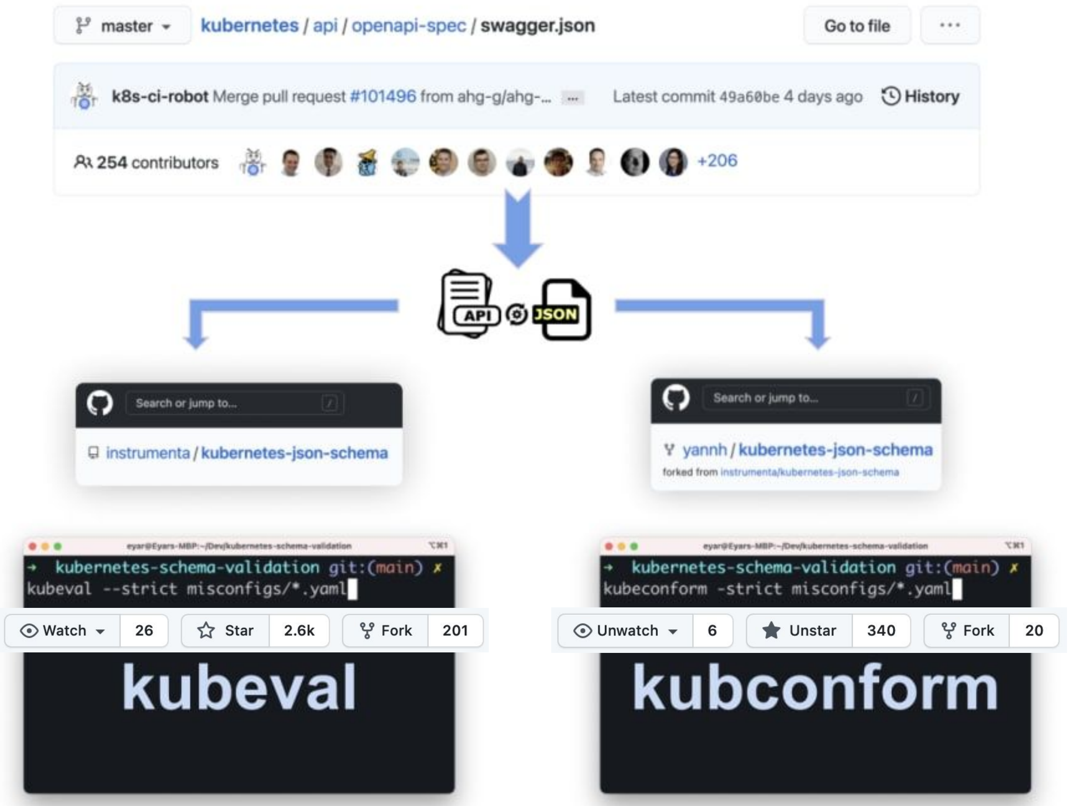

master / kubernetes / api / openapi-spec / swagger.json

k8s-ci-robot Merge pull request #101496 from ahg-g/ahg-... Latest commit 49a60be 4 days ago History

254 contributors +206

instrumenta / kubernetes-json-schema

yannh / kubernetes-json-schema
forked from instrumenta/kubernetes-json-schema

```
eyar@Eyars-MBP:~/Dev/kubernetes-schema-validation
→  kubernetes-schema-validation git:(main) ✗
kubeval --strict misconfigs/*.yaml
```

Watch 26 | Star 2.6k | Fork 201

kubeval

```
eyar@Eyars-MBP:~/Dev/kubernetes-schema-validation
→  kubernetes-schema-validation git:(main) ✗
kubeconform -strict misconfigs/*.yaml
```

Unwatch 6 | Unstar 340 | Fork 20

kubconform

kubectl vs. open-source

Let’s QA - take #1

invalid-labels-value.yaml

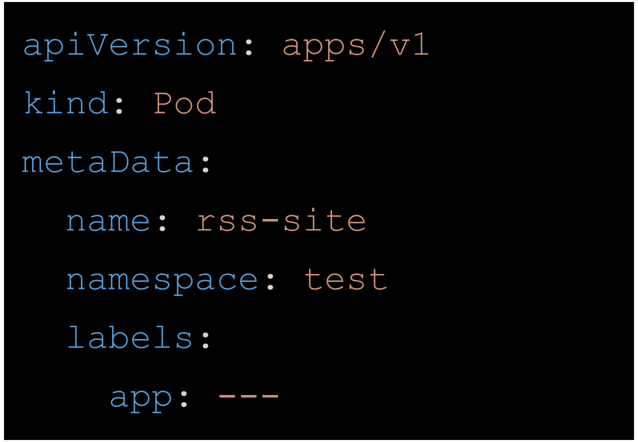

```yaml
apiVersion: apps/v1
kind: Pod
metaData:
  name: rss-site
  namespace: test
  labels:
    app: ---
```

[Catching] kubectl --dry-run=server

[Not catching] kubeval / kubeconform

The Pod "invalid-labels-value" is invalid: metadata.labels: Invalid value: "---": a valid label must be an empty string or consist of alphanumeric characters, '-', '' or '.', and must start and end with an alphanumeric character (e.g. 'MyValue', or 'my_value', or '12345', regex used for validation is '(([A-Za-z0-9][-A-Za-z0-9.]*)?[A-Za-z0-9])?')

Let’s QA - take #2

missing-image.yaml

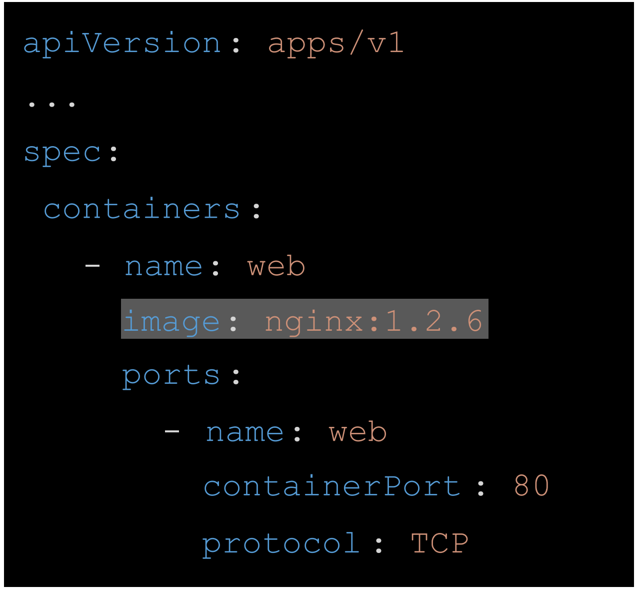

```
apiVersion: apps/v1
...
spec:
 containers:
    - name: web
      image: nginx:1.2.6
      ports:
        - name: web
          containerPort : 80
          protocol : TCP
```

[Catching] kubectl --dry-run=server

[Not catching] kubeval / kubeconform

Winner?

A connection to your cluster is allowed?

kubectl --dry-run=server

A connection to your cluster is *NOT* allowed?

Kubeval Kubeconform

kubeval vs. kubeconform

Kubernetes versions support

👎 Kubeval - instrumenta/kubernetes-json-schema

(last commit: 133f848 on April 29, 2020)

- v1.5.0 - v1.18.1

👍 Kubeconform - yannh/kubernetes-json-schema

(last commit: 14652b0on Nov 19, 2021)

- v1.15.0 - v1.22.4

CRDs support

👎 Kubeval:

- “This means it's **not possible to validate resources using CRDs**. Currently you need to pass a flag to ignore missing schemas…”

👍 Kubeconform:

- “Overriding schemas location - CRD and Openshift support”

Community

Kubeval:

👍 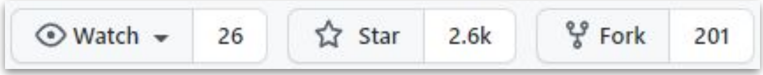

👎 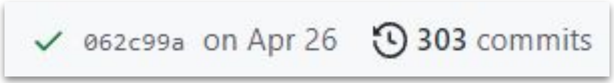
✓ 062c99a on Apr 26 303 commits

Kubeconform:

😐 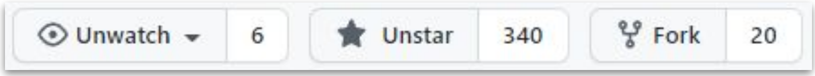

👍 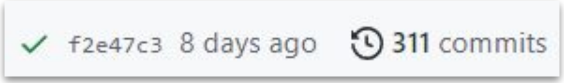
✓ f2e47c3 8 days ago 311 commits

Winner?

kubeconform!

Strategies for validating Kubernetes schema

- ⬅️ Shift-left
- 🚔 Fill the gap
- 👯 Alternatives
- 💸 Buy vs. build

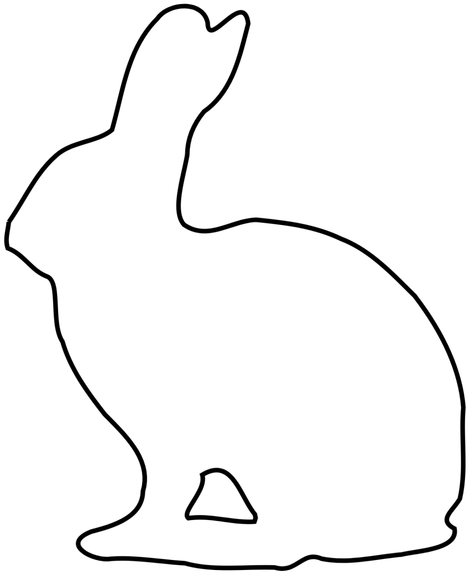

Strategies for validating Kubernetes schema

⬅️ Shift-left

When possible, the best setup is if you can run kubectl --dry-run=server on every code change, but you probably can’t do it because you can’t allow every developer or CI machine in your organization to have a connection to your cluster. So, the second-best effort is to run kubeconform.

Strategies for validating Kubernetes schema

🚔 Fill the gap

Because kubeconform doesn't cover all common misconfigurations, it's recommended to run it with a policy enforcement tool on every code change to fill the coverage gap.

Strategies for validating Kubernetes schema

👯 Alternatives

Another option for using kubectl dry-run in server mode, without having a connection to your Kubernetes environment, is to run minikube + kubectl --dry-run=server. The downside of this hack is that it's also required to set up the minikube cluster like prod (same volumes, namespace, etc.) or you'll encounter errors when trying to validate your Kubernetes manifests.

Strategies for validating Kubernetes schema

💸 Buy vs. build

If you enjoy the engineering overhead, then kubeconform + conftest is a great combination of tools to get good coverage. Alternatively, there are tools that can provide you with an out-of-the-box experience to help you save time and resources, such as Datree (whose schema validation is powered by kubeconform).

Thank you

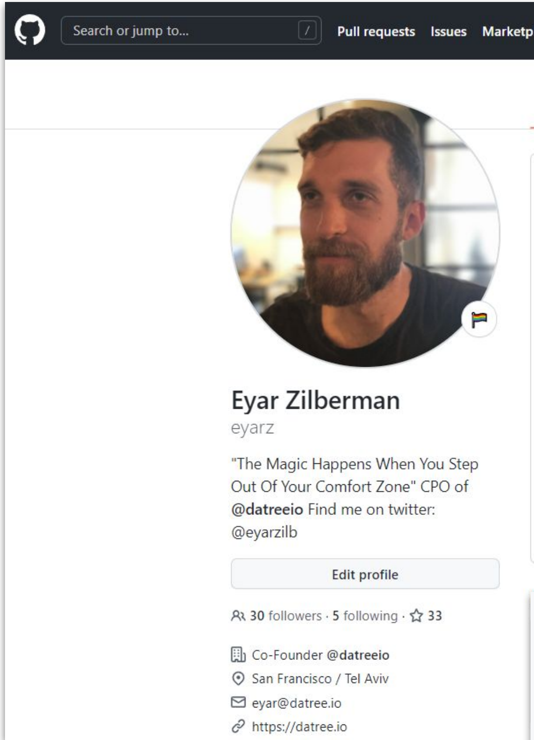

Eyar Zilberman

eyarz

"The Magic Happens When You Step Out Of Your Comfort Zone" CPO of @datreeio Find me on twitter: @eyarzilb

Edit profile

30 followers · 5 following · ☆ 33

Co-Founder @datreeio

San Francisco / Tel Aviv

eyar@datree.io

https://datree.io

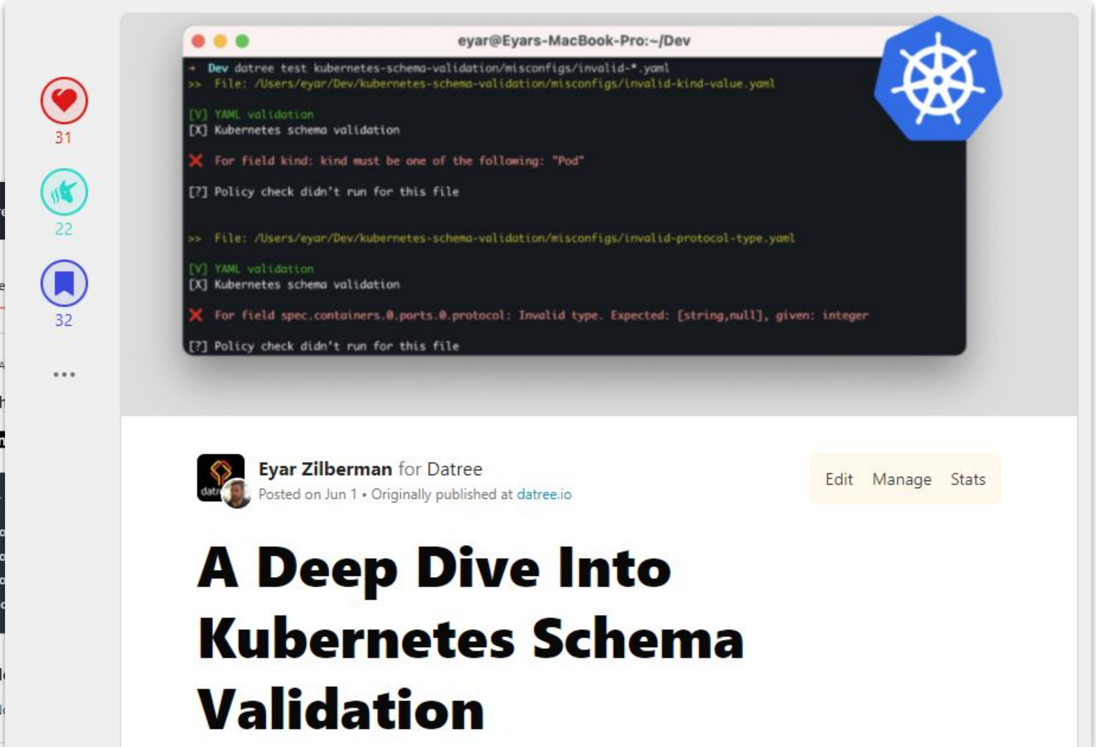

eyar@Eyars-MacBook-Pro:~/Dev

```
→ Dev datree test kubernetes-schema-validation/misconfigs/invalid-*.yaml
>>  File: /Users/eyar/Dev/kubernetes-schema-validation/misconfigs/invalid-kind-value.yaml

[V] YAML validation
[X] Kubernetes schema validation

❌ For field kind: kind must be one of the following: "Pod"

[?] Policy check didn't run for this file


>>  File: /Users/eyar/Dev/kubernetes-schema-validation/misconfigs/invalid-protocol-type.yaml

[V] YAML validation
[X] Kubernetes schema validation

❌ For field spec.containers.0.ports.0.protocol: Invalid type. Expected: [string,null], given: integer

[?] Policy check didn't run for this file
```

Eyar Zilberman for Datree

Posted on Jun 1 • Originally published at datree.io

A Deep Dive Into Kubernetes Schema Validation

eyarz / pink-bunny-ears Public

datree.io