

# Scaling Observability

Harnessing the Power of Kubernetes Operators



 [chrismuellner](#)

 [christoph-muellner-at](#)




# Christoph Müllner

Sr. Software Engineer & Product Owner at Dynatrace

# Agenda

---

- Background
- Operator
- Webhook
- CSI Driver





Cloud done right.

Hybrid cloud and Kubernetes analytics

Hybrid cloud distributed tracing

Attack blocking and protection

Mobile, web browser and API

Real-time business insights

Closed-loop remediation

Ecosystem integrations

Log and event management

Automatic code-level root-cause and profiling

Vulnerability runtime analytics

Feature adoption analysis

Impact and conversion

Quality gate, service level and delivery

Custom solutions

Automatic enterprise-grade observability

Front- & back-end availability and performance

Risk-based remediation

Visual session replay

BizDevOps integration and automation

DevOps, SRE lifecycle

API programmability

Infrastructure Monitoring

Applications & Microservices

Application Security

Digital Experience

Business Analytics

Cloud Automation

Dynatrace Hub

dynatrace Software Intelligence Platform

OneAgent®

PurePath®

Smartscape®

Grail™

Davis® AI

Traces Metrics Logs

+

Topology Behaviour Code Metadata Network

+

API OpenTelemetry keptn

600+

Supported technologies

Kubernetes

OpenShift

AWS

Azure

GCP

Tanzu

Enterprise

Hybrid cloud

Automatic and intelligent observability

Broadest multicloud and technology support

# OneAgent



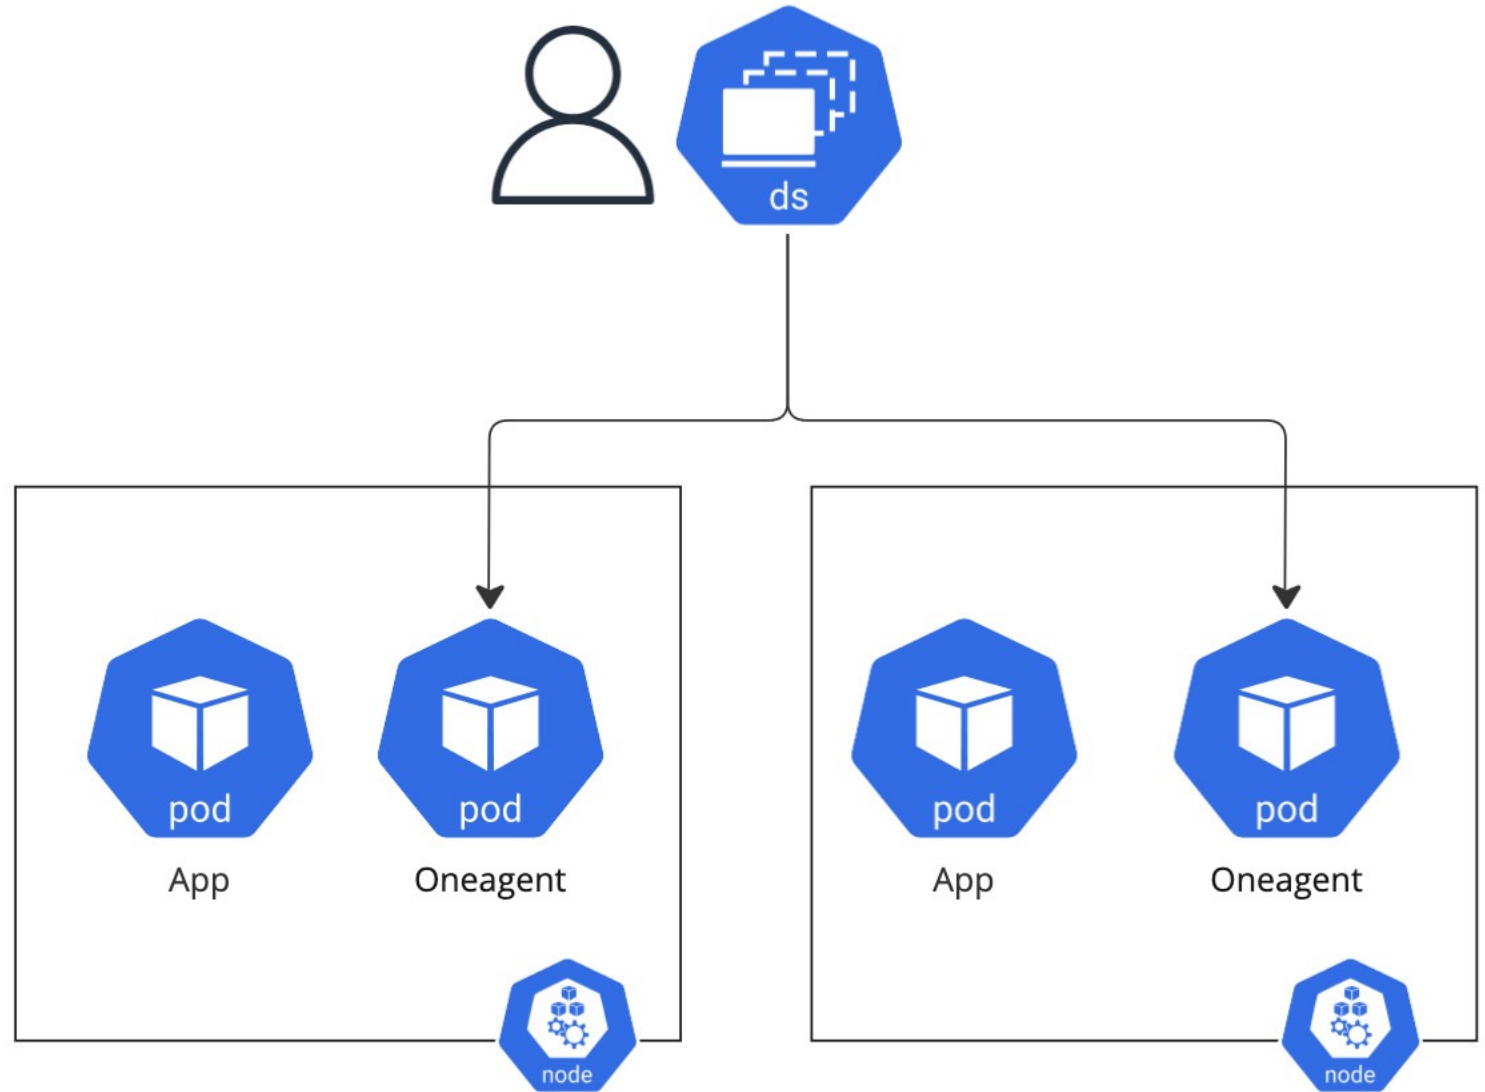
# Concept

---

- Inject applications without manual configuration
- Automatically attach necessary metadata/binaries
  - LD\_PRELOAD environment variable
  - binaries via mount points
- *Needs to be present before process startup!*

# Deploy OneAgent

- Put container in pod
- Schedule pod on every node via daemonset




# Problems

- Unmanaged manifest

```
1 apiVersion: apps/v1
2 kind: DaemonSet
3 metadata:
4   labels:
5     app.kubernetes.io/component: fullstack
6     app.kubernetes.io/name: oneagent
7     app.kubernetes.io/version: 1.277.0.20230914-094321
8   name: oneagent
9   namespace: dynatrace
10 spec:
11   revisionHistoryLimit: 10
12   selector:
13     matchLabels:
14       app.kubernetes.io/name: oneagent
15   template:
16     metadata:
17       annotations:
18         container.apparmor.security.beta.kubernetes.io/dynatrace-oneagent: unconfined
19       labels:
20         app.kubernetes.io/component: fullstack
21         app.kubernetes.io/name: oneagent
22         app.kubernetes.io/version: 1.277.0.20230914-094321
23     spec:
24       affinity:
25         nodeAffinity:
26           requiredDuringSchedulingIgnoredDuringExecution:
27             nodeSelectorTerms:
28               - matchExpressions:
29                 - key: kubernetes.io/arch
30                   operator: In
31                   values:
32                     - amd64
```

- Manual updates





# How to best deploy it on k8s?

---

## Helm

- simple deployment
- single rollout
- basic validation



- *can be used for operator deployment!*

## Operator

- lifecycle management
- fine-grained validation



# Operator



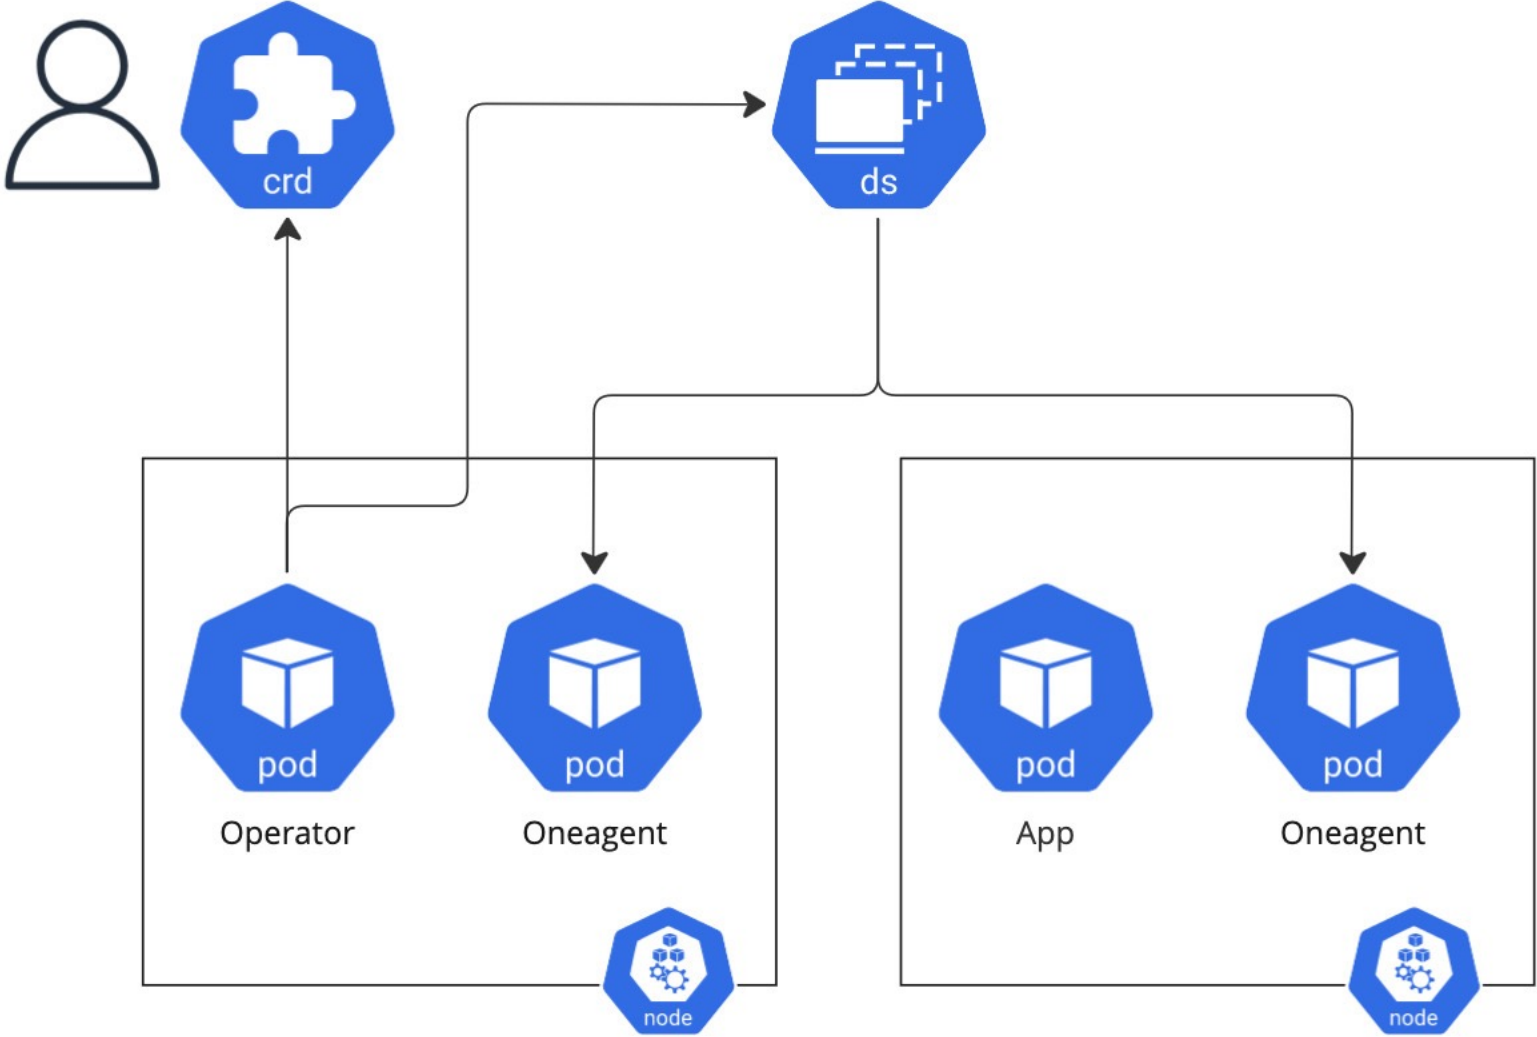
# Concept

---

- Configure via Custom Resource
- Reconcile resources

```
1 apiVersion: dynatrace.com/v1beta1
2 kind: DynaKube
3 metadata:
4   name: dynakube
5 spec:
6   apiUrl: TENANT.dynatracelabs.com/api
7
8   oneAgent:
9     cloudNativeFullStack: {}
```

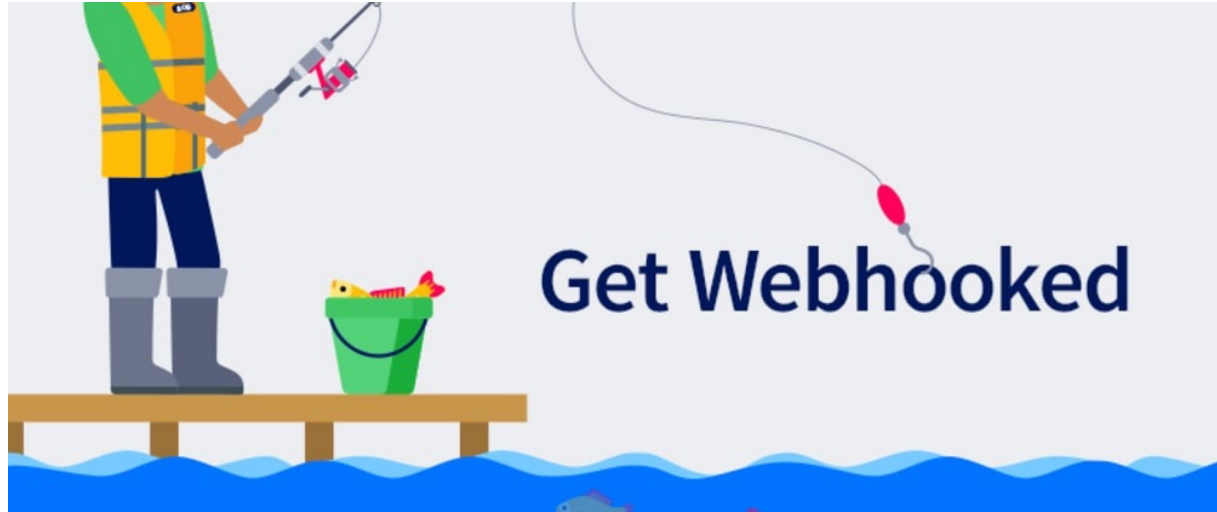
# Operator deployment



# Problems

---

- Container runtime injection
  - hidden from Cluster Operators
- Race Condition
  - container needs to be ready on node to inject pods
- No Multi-tenancy
  - single Oneagent per node



# Webhook

# Concept

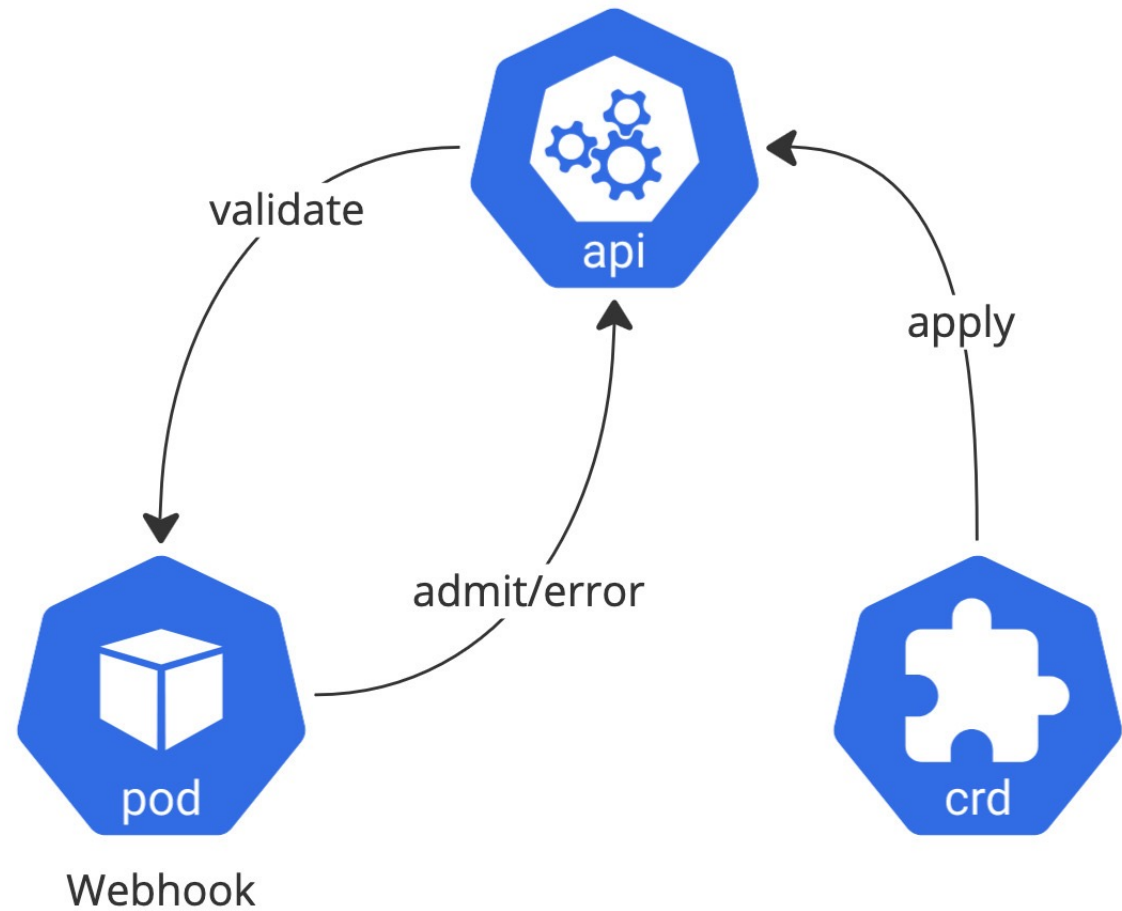
---

- Event-based
- Different types
  - Validating
  - Mutating
  - Conversion



# Validation webhook

- Validate configuration
  - Early feedback
  - Clear error message





# Validation webhook

---



```
1 apiVersion: dynatrace.com/v1beta1
2 kind: DynaKube
3 metadata:
4   name: invalid-dynakube
5 spec:
6   apiUrl: ""
```

Error from server (Forbidden): error when creating "STDIN": admission webhook "webhook.dynatrace.com" denied the request:

2 error(s) found in the Dynakube

1. The DynaKube's specification is missing the API URL or still has the example value set.

Make sure you correctly specify the URL in your custom resource.

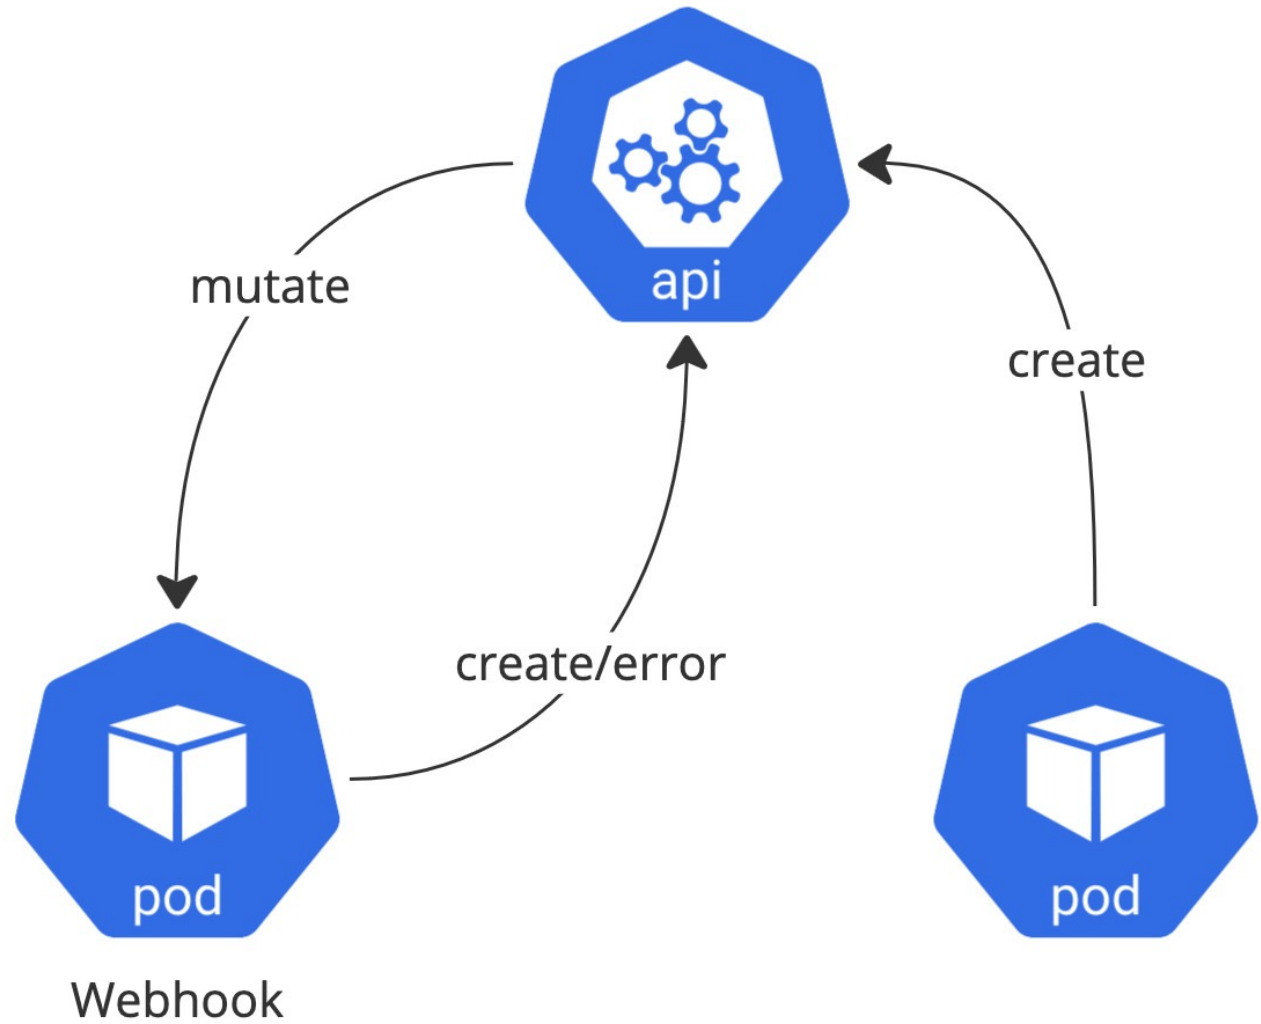
2. The DynaKube's specification has an invalid API URL value set.

Make sure you correctly specify the URL in your custom resource (including the /api postfix).




# Mutate pods

- intercept create/update/... events
- mutate/error pod

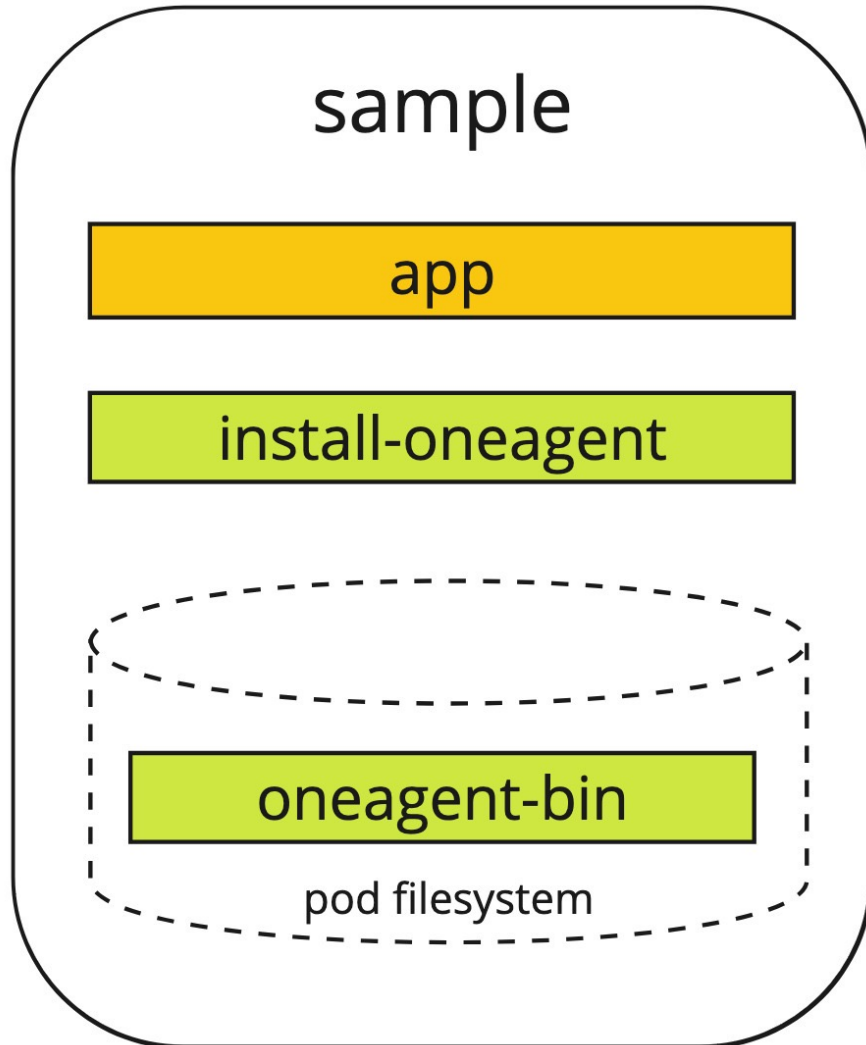


# Mutate pods



```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4   name: sample
5 spec:
6   containers:
7     - image: fancyapp
8       name: app
9       ports:
10      - containerPort: 80
11        protocol: TCP
```

# Mutate pods



```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4   name: sample
5 spec:
6   containers:
7     - name: app
8       image: fancyapp
9       ports:
10      - containerPort: 80
11      env:
12      - name: LD_PRELOAD
13        value: /opt/dynatrace/oneagent-paas/agent/lib64/liboneagentproc.so
14      volumeMounts:
15      - mountPath: /opt/dynatrace/oneagent-paas
16        name: oneagent-bin
17   initContainers:
18     image: docker.io/dynatrace/dynatrace-operator
19     name: install-oneagent
20     volumeMounts:
21     - mountPath: /mnt/bin
22       name: oneagent-bin
23   volumes:
24     - emptyDir: {}
25       name: oneagent-bin
```

# Configuration

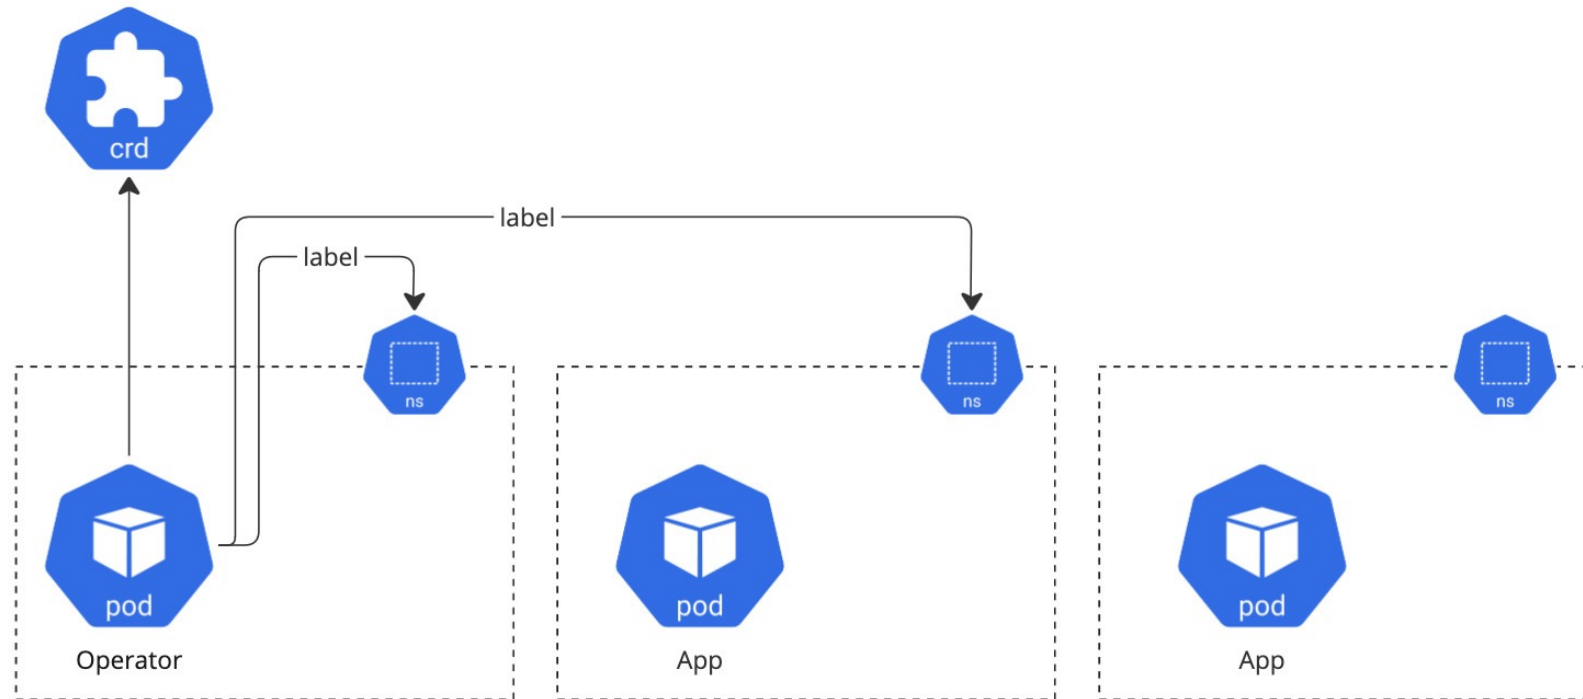
---

- cluster-wide by default
  - exceptions apply
- namespace selector
  - *manually label every namespace?*

```
1 apiVersion: admissionregistration.k8s.io/v1
2 kind: MutatingWebhookConfiguration
3 metadata:
4   name: dynatrace-webhook
5 webhooks:
6   - name: webhook.pod.dynatrace.com
7   ...
8     namespaceSelector:
9       matchExpressions:
10        - key: dynakube.internal.dynatrace.com/instance
11          operator: Exists
```

# Configuration

- namespace selector
  - webhook matches no namespace by default
  - namespace label managed by Operator according to config



# Fun with other webhooks

- alphabetical order
  - unknown by webhooks
- reinvocation policy
  - reinvoked if pod is modified by other webhook
  - annotations to optimise reinvoked requests



# Resilience

---

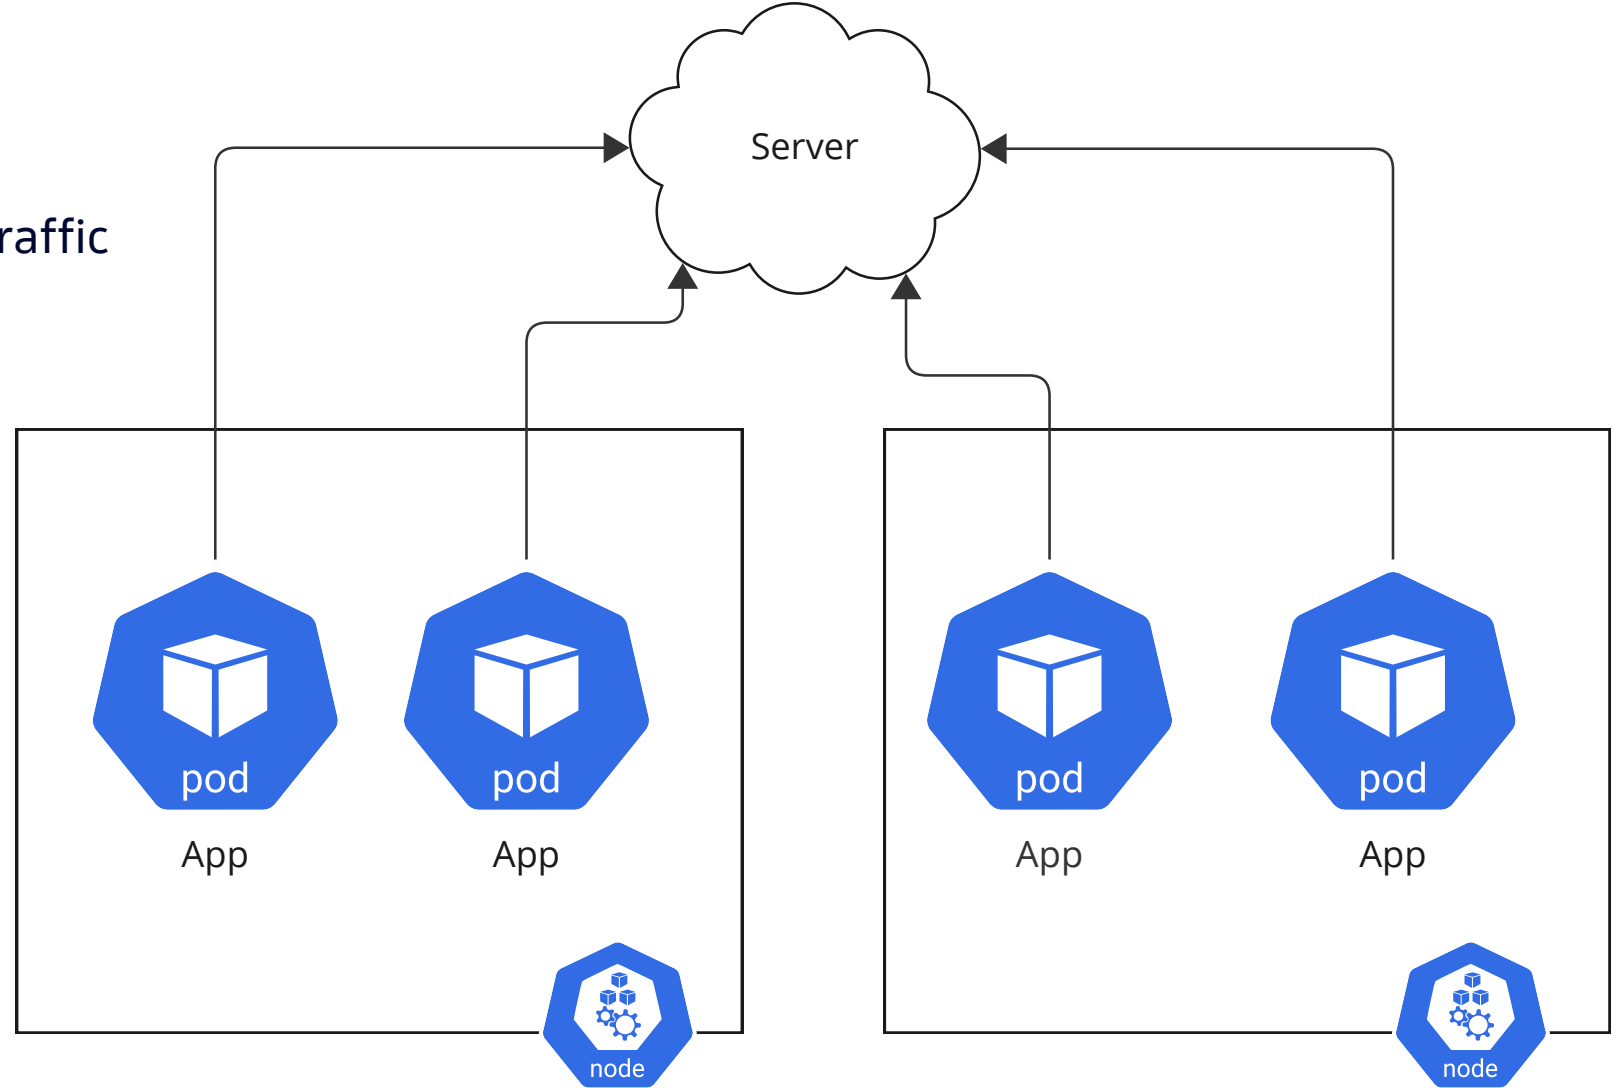
- Webhook
  - topology spread constraints
  - failure policy & timeout
- Init container
  - custom failure policy





# Problems

- Download per pod
  - Increased storage & network traffic
  - Increased startup time





# CSI Driver

# Concept

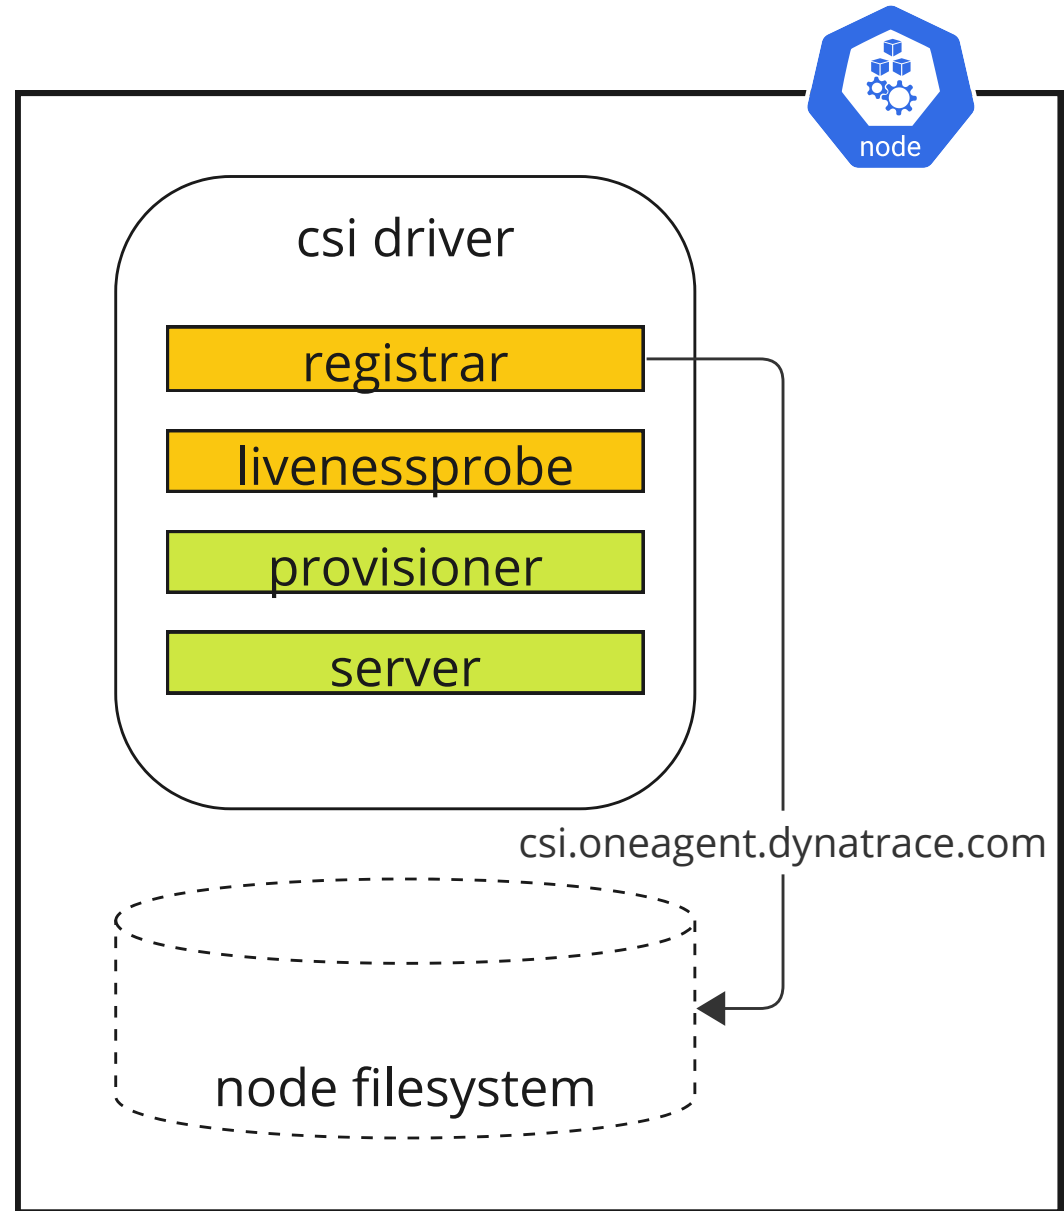
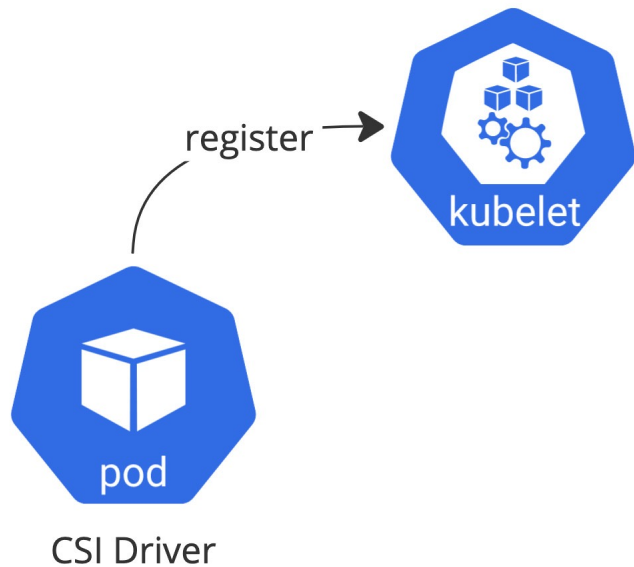
---

- Container Storage Interface
- CSI object to register

```
1 apiVersion: storage.k8s.io/v1
2 kind: CSIDriver
3 metadata:
4   name: csi.oneagent.dynatrace.com
5 spec:
6   attachRequired: false
7   podInfoOnMount: true
8   volumeLifecycleModes:
9     - Ephemeral
```

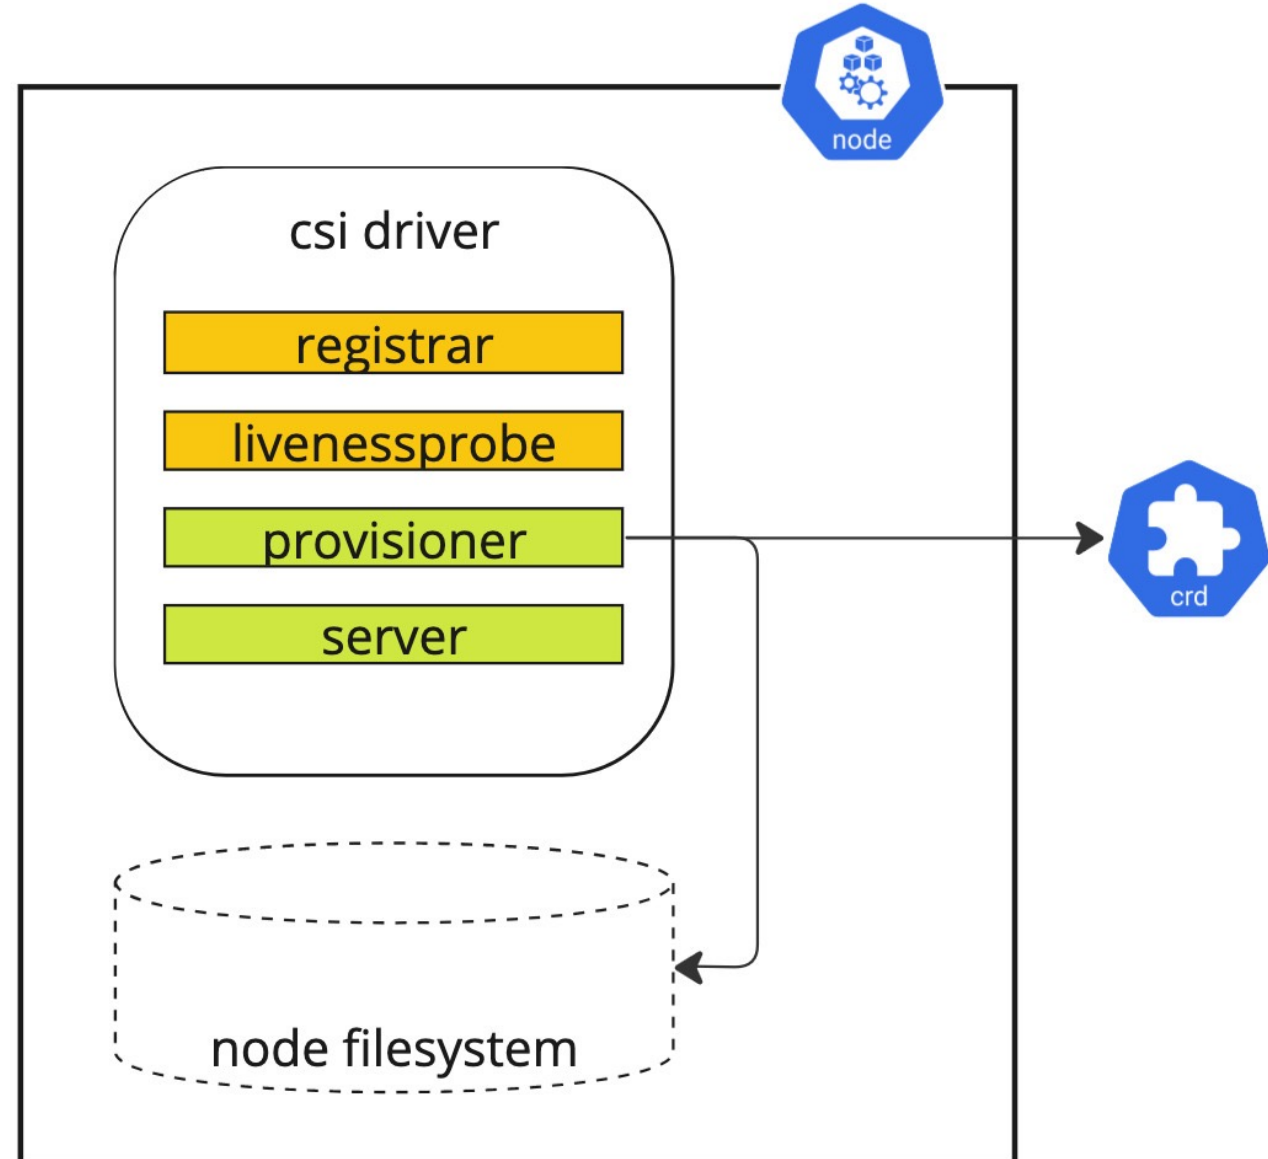
# Concept

- CSI sidecars
  - registrar
  - liveness probe
- liveness probe

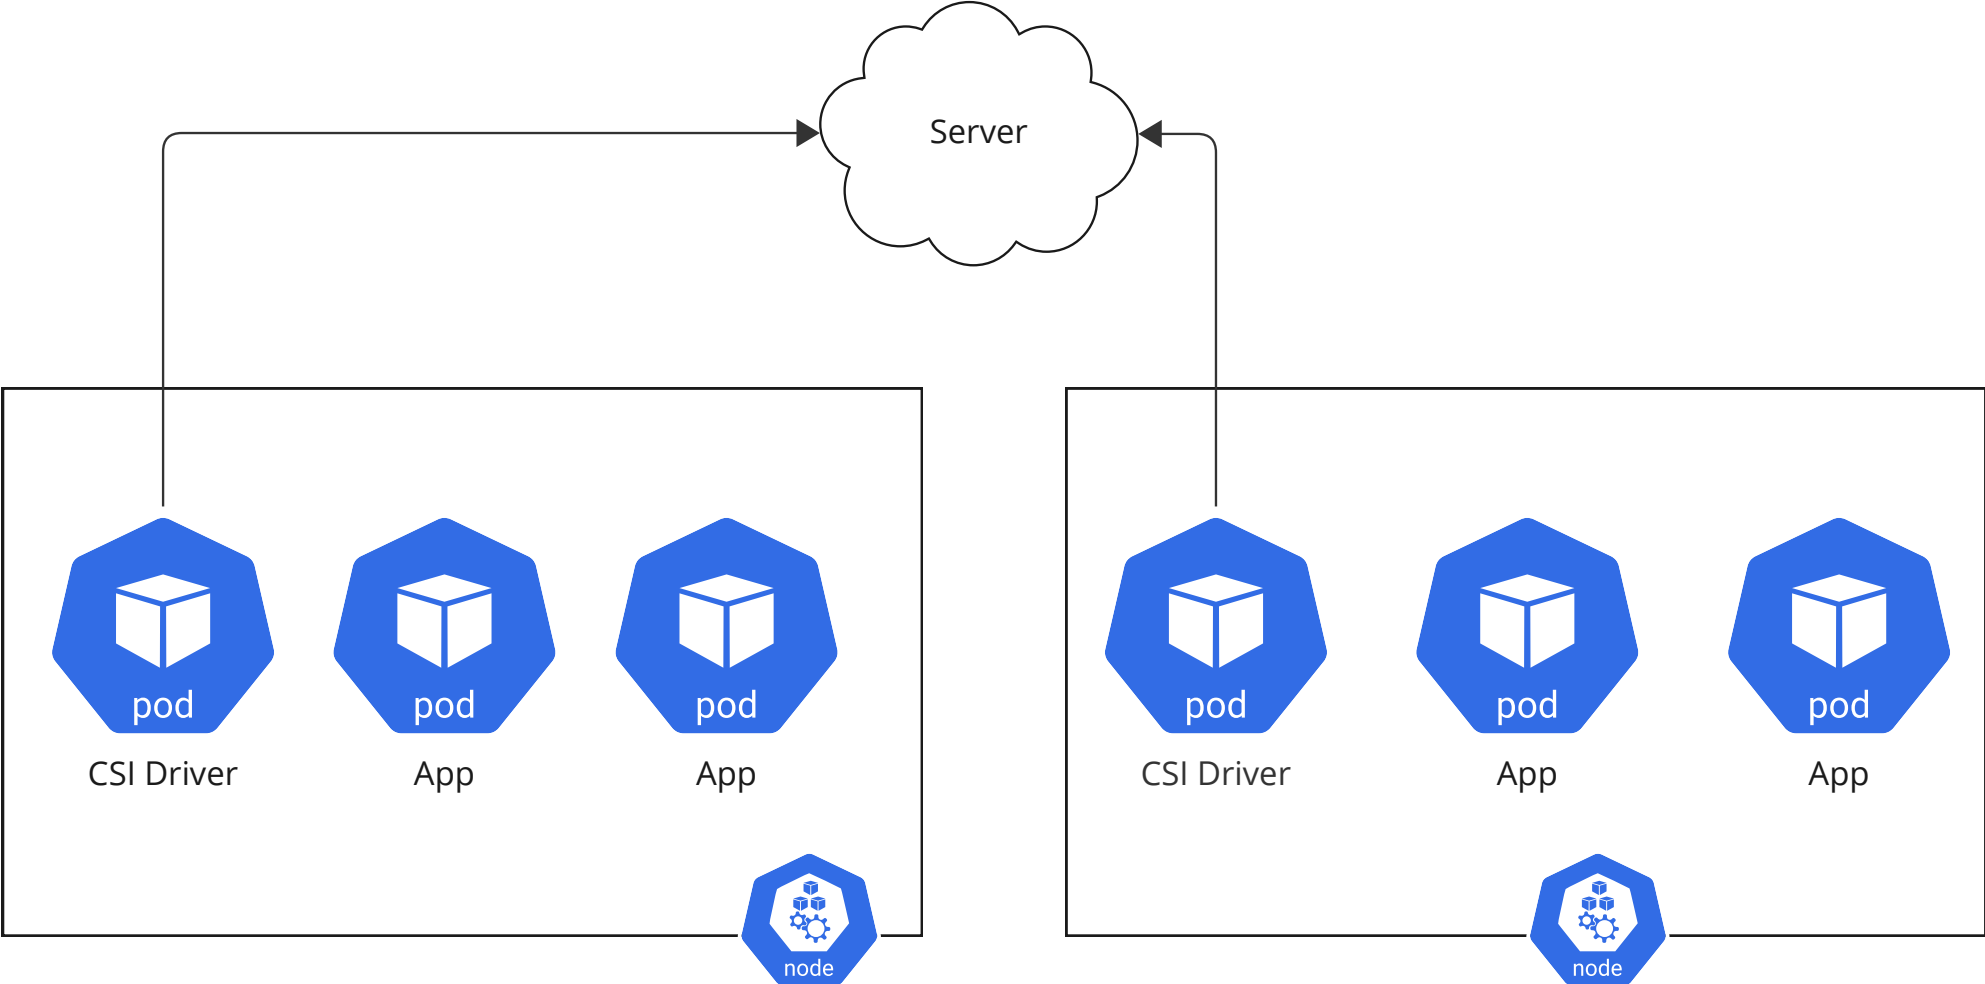


# Binaries cached on node

- Watch Custom Resources
- Download required agent versions

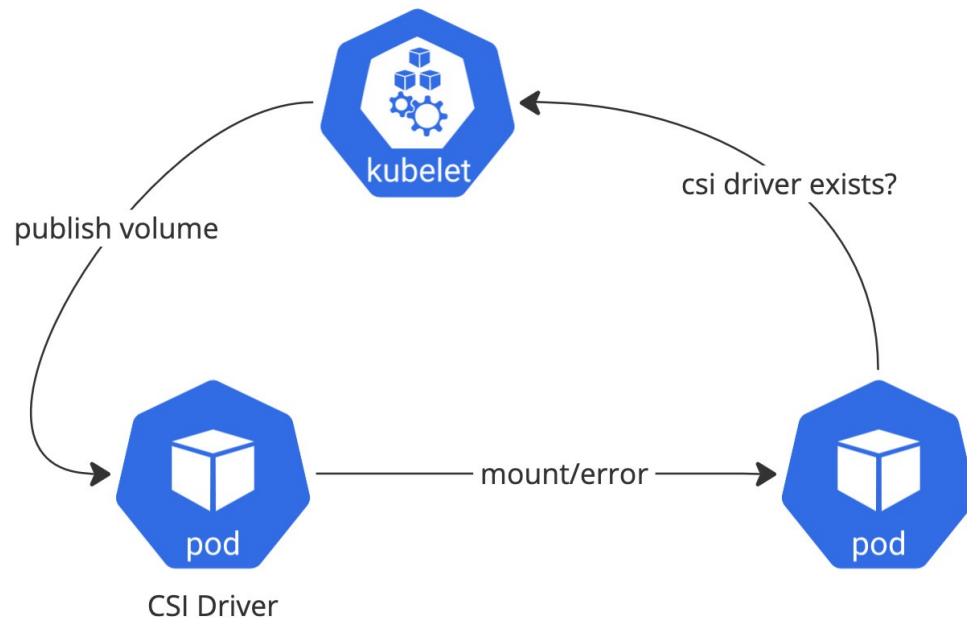


# Binaries cached on node



# Custom mount logic

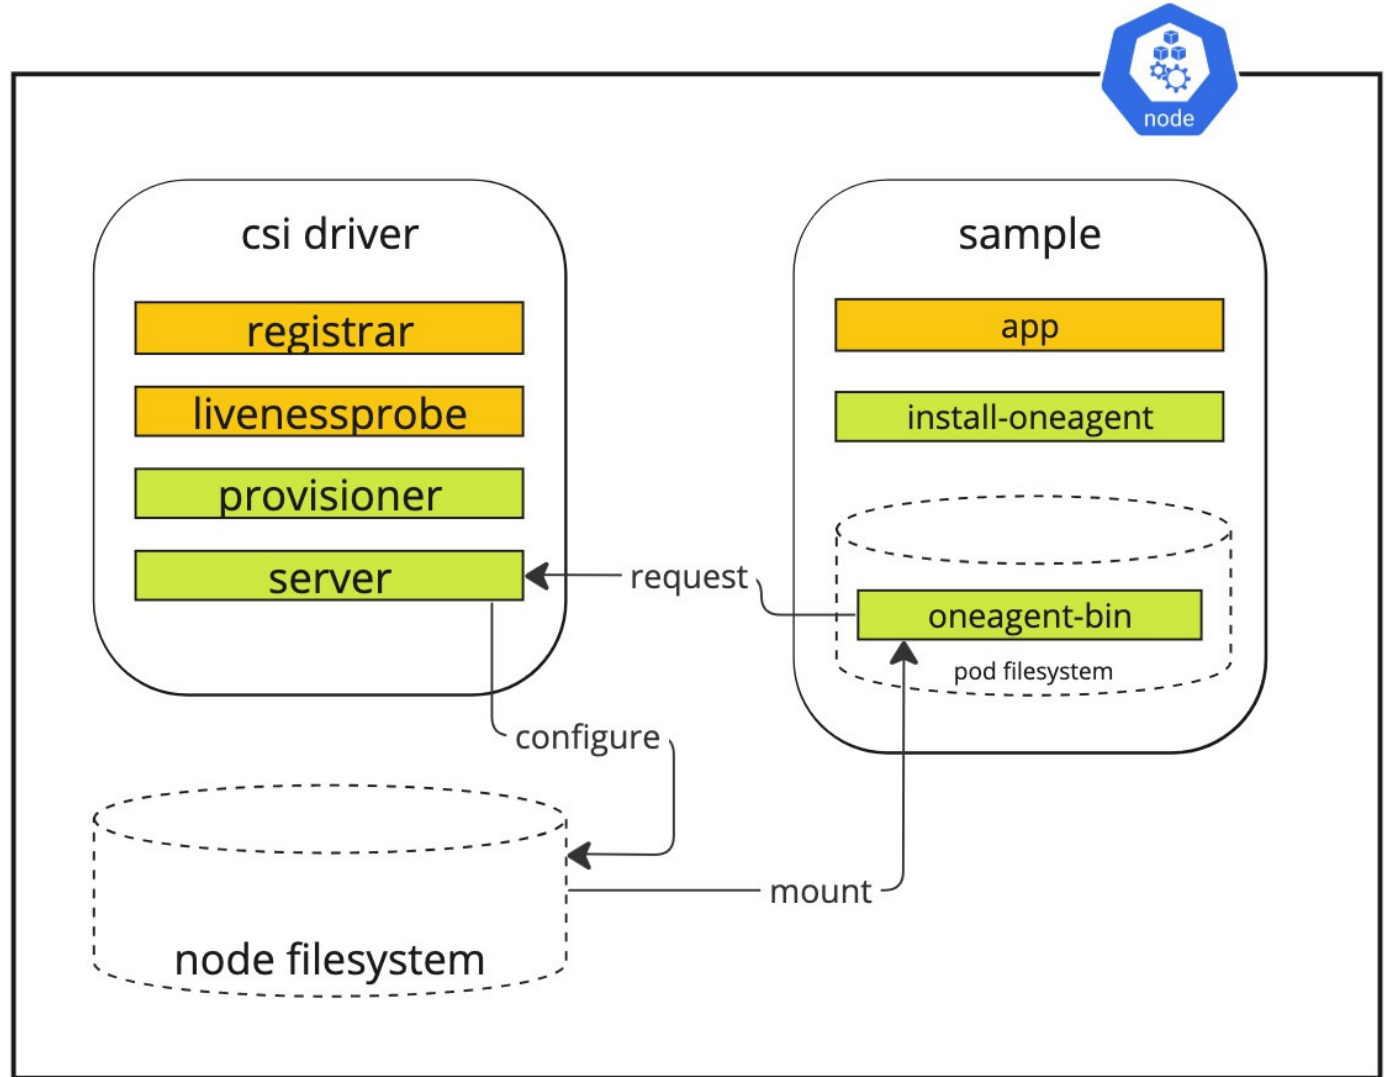
- Volume handled by CSI driver



```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4   name: sample
5 spec:
6   containers:
7     - name: app
8       image: fancyimage
9       ports:
10        - containerPort: 80
11       env:
12        - name: LD_PRELOAD
13          value: /opt/dynatrace/oneagent-paas/agent/lib64/liboneagentproc.so
14       volumeMounts:
15        - mountPath: /opt/dynatrace/oneagent-paas
16          name: oneagent-bin
17   initContainers:
18     image: docker.io/dynatrace/dynatrace-operator:snapshot
19     name: install-oneagent
20     volumeMounts:
21      - mountPath: /mnt/bin
22        name: oneagent-bin
23   volumes:
24     - csi:
25       driver: csi.oneagent.dynatrace.com
26       readOnly: false
27       volumeAttributes:
28        dynakube: dynakube
29        mode: app
30       name: oneagent-bin
```

# Custom mount logic

- Mount correct agent version
  - binaries linked via overlayFS
  - configuration adapted per pod





# Resilience

---

- Split containers
  - independantly failing containers
  - server can always handle mount requests
- Custom timeout
  - count failed mount attempts
  - mount empty volume



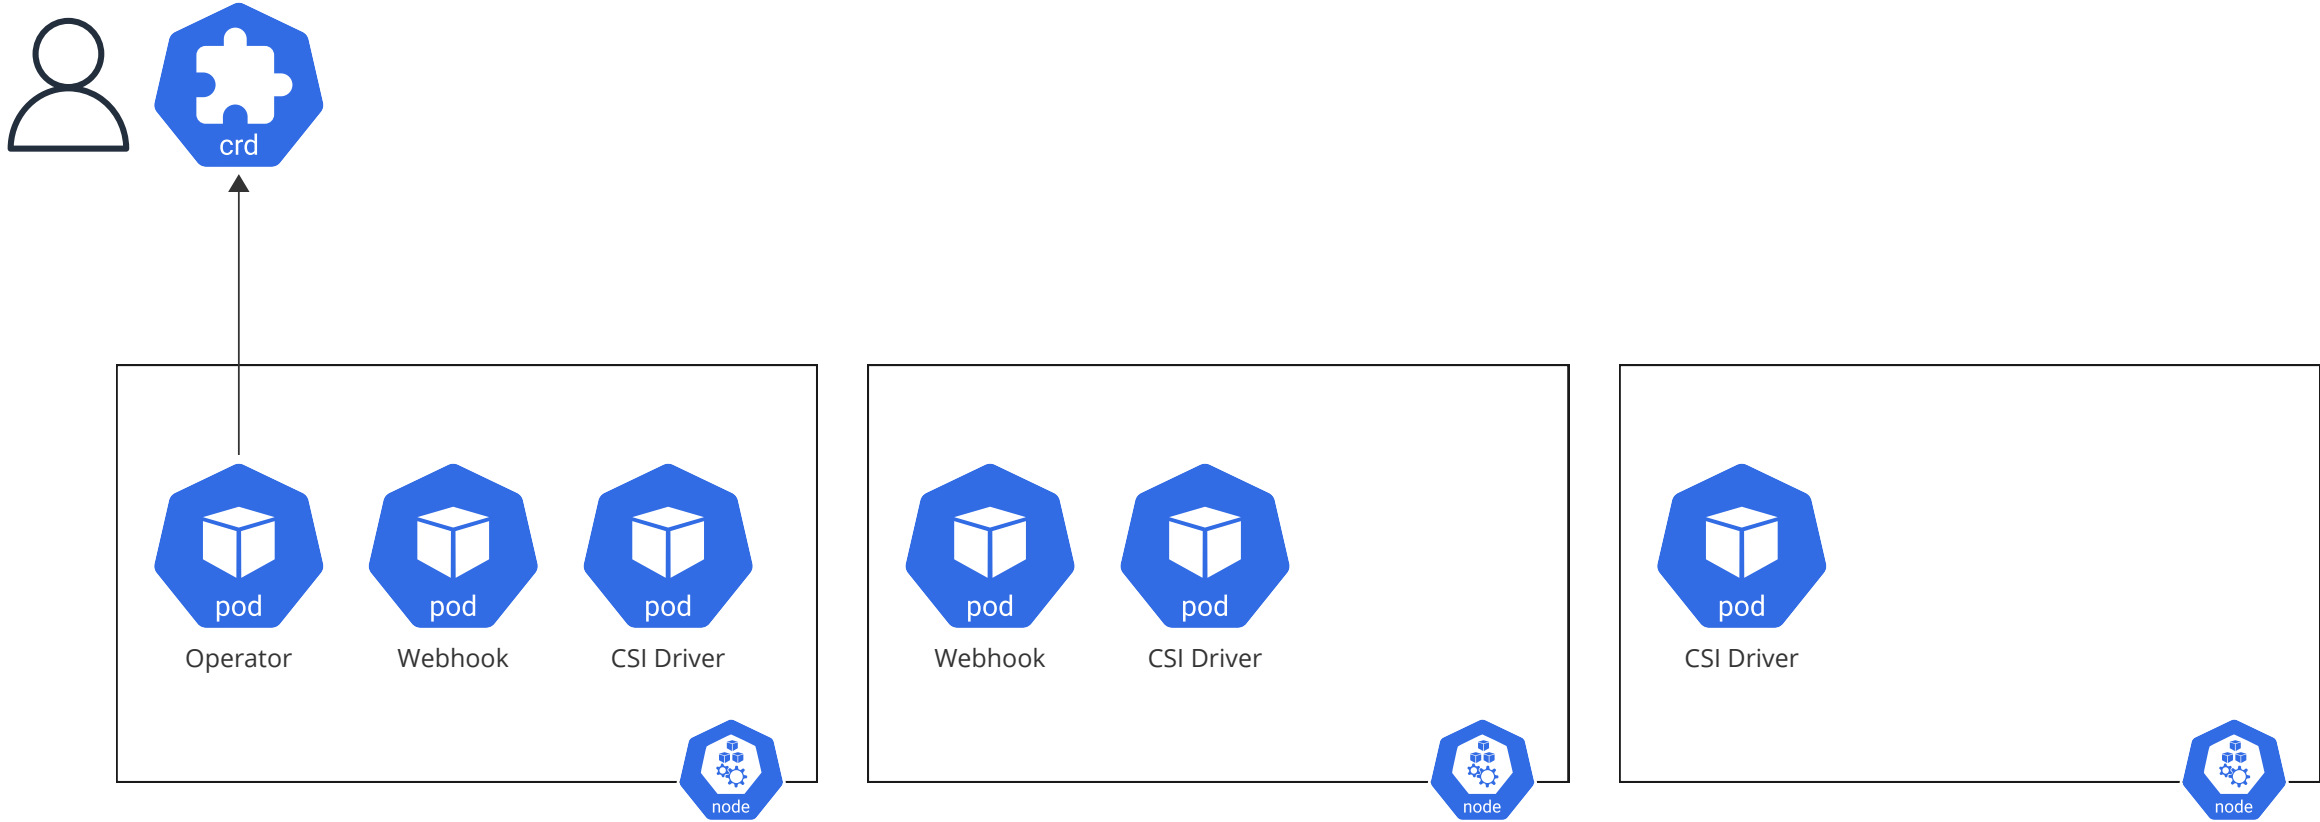
# Problems

---

- requires privileged container
- Security Context Constraints in Openshift
  - Openshift <4.13 does not allow CSI volumes by default
- Need to garbage collect old agents on node filesystem



# Overview



# Thanks!

## Scaling Observability

### Dynatrace/**dynatrace-operator**



Automate Kubernetes observability with Dynatrace

 36  
Contributors

 2  
Issues

 131  
Stars

 129  
Forks

