

Kubelet Deep Dive: Writing a Kubelet in Rust

Kevin Flansburg



About Me



KubeCon




CloudNativeCon

North America 2020

Virtual

- Rust - 3 years
 - Microservices
 - DevOps
- Kubernetes - 1 year
- Krustlet maintainer



Krustlet Project



KubeCon



CloudNativeCon

North America 2020

Virtual

- **Kubernetes Rust Kubelet**
- [Deis Labs](#)
- [GitHub Repository](#)

- `kubelet` - crate for building Kubelets
- Kubelet Implementations:
 - waSCC
 - WASI
 - Linux Containers via CRI

Kubernetes Architecture



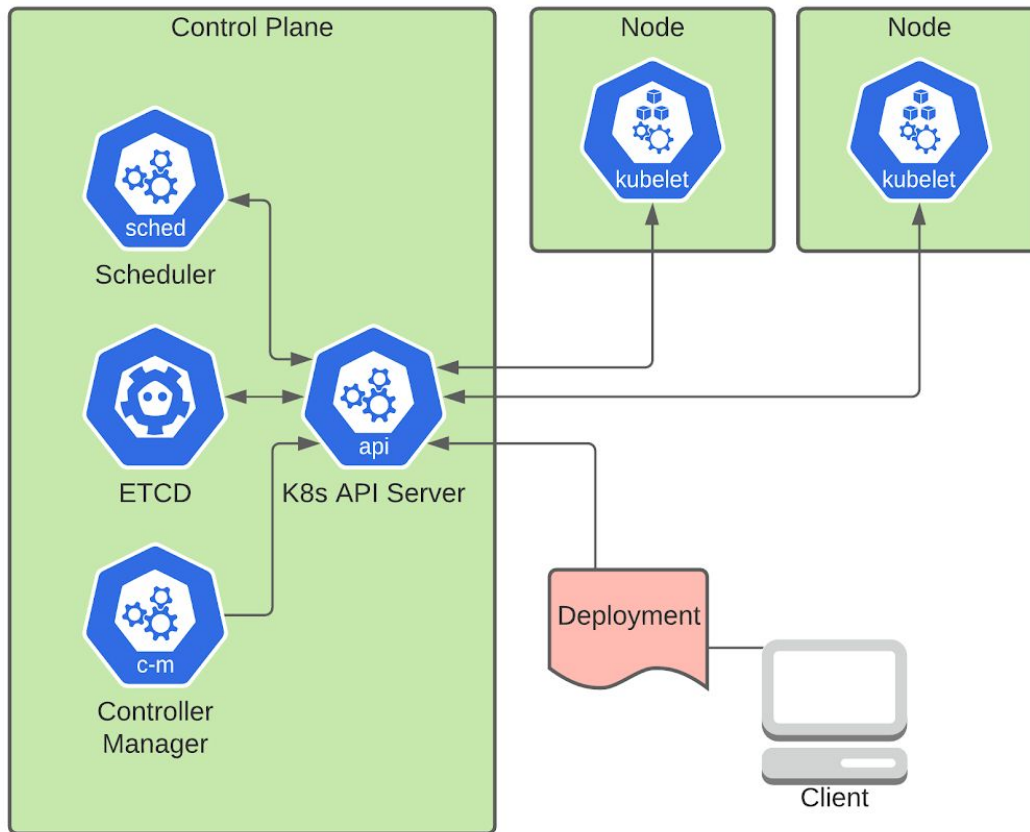
KubeCon



CloudNativeCon

North America 2020

Virtual



Controller Pattern



KubeCon




CloudNativeCon

North America 2020

Virtual

- (Mostly) immutable resource objects.
- Monitor for changes to objects of a particular resource.
 - Add / Modify / Delete
 - “Informer” pattern
- Drive cluster state to match this desired state.



Aside: Operators

- Controller with:
 - Custom Resource Definition
 - Application / domain specific



Kubelet: Franken-controller



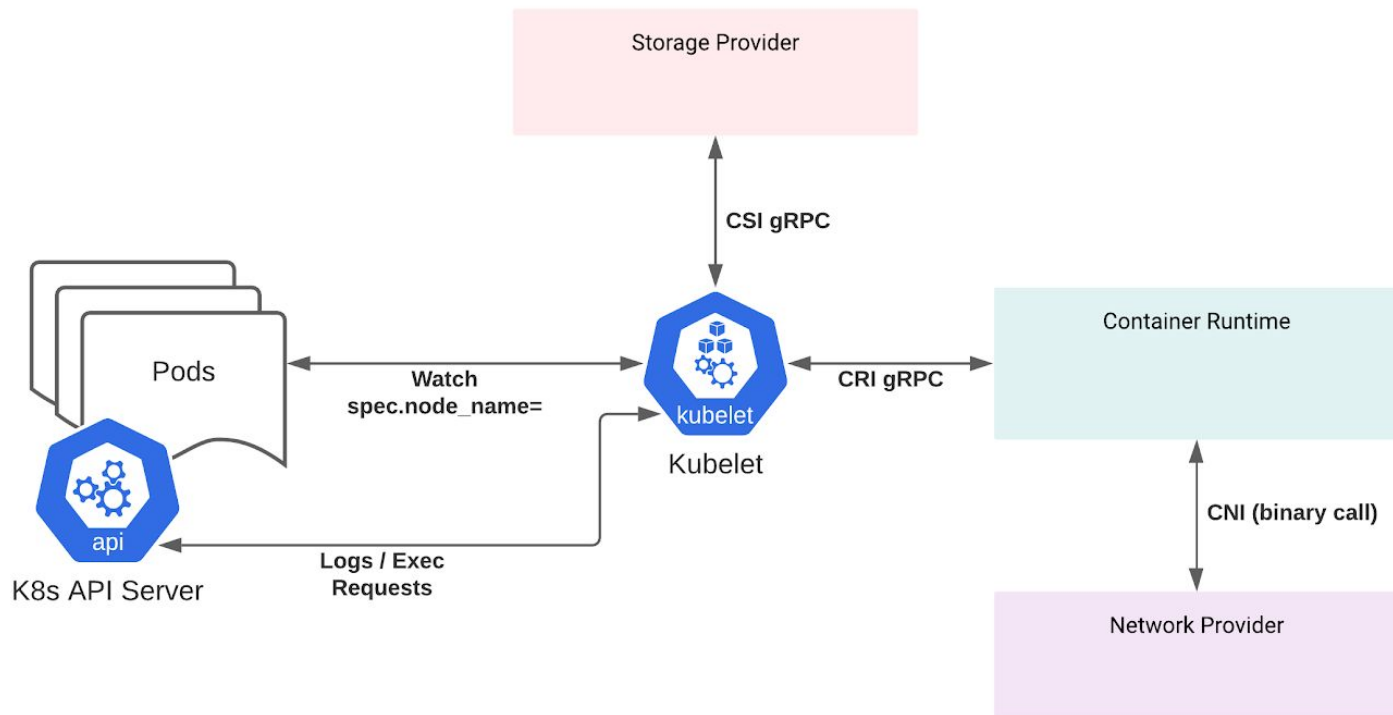
KubeCon



CloudNativeCon

North America 2020

Virtual



Rust for Distributed Apps



KubeCon



CloudNativeCon

North America 2020

Virtual

- Performance
- Strongly typed
- Ownership - “fearless concurrency”
- `async/await`
- Error handling

```
let result = failable_fn();
match result {
    Ok(value) => ...,
    Err(e) => ...,
}
```

Or simply:

```
let result = failable_fn()?;
```



```
warning: variable does not need to be mutable
--> src/main.rs:1:14
1 | fn increment(mut x: u64) -> u64 {
  |               ^
  |               help: remove this `mut`
= note: `#[warn(unused_mut)]` on by default
```


Rust Ecosystem



KubeCon



CloudNativeCon

North America 2020

Virtual

- Many fantastic crates
 - [serde](#)
 - [tracing](#)
 - [prost](#) / [tonic](#)
- Documentation
- Dependency management
- Great community

```
src/main.rs

use serde::{Serialize, Deserialize};

#[derive(Serialize, Deserialize, Debug)]
struct Point {
    x: i32,
    y: i32,
}

fn main() {
    let point = Point { x: 1, y: 2 };

    let serialized = serde_json::to_string(&point).unwrap();
    println!("serialized = {}", serialized);

    let deserialized: Point = serde_json::from_str(&serialized).unwrap();
    println!("deserialized = {:?}", deserialized);
}
```

Run

Serde Demo in Runnable Documentation Example

Useful Kubernetes Crates



KubeCon



CloudNativeCon

North America 2020

Virtual

- [k8s-openapi](#) - Rust types for Kubernetes API resources.
- [kube](#) - Kubernetes client.
- [k8s-cri](#) - gRPC client for Container Runtime Interface (CRI)
- [k8s-csi](#) - gRPC client for Container Storage Interface (CSI)

Kubelet Control Loop



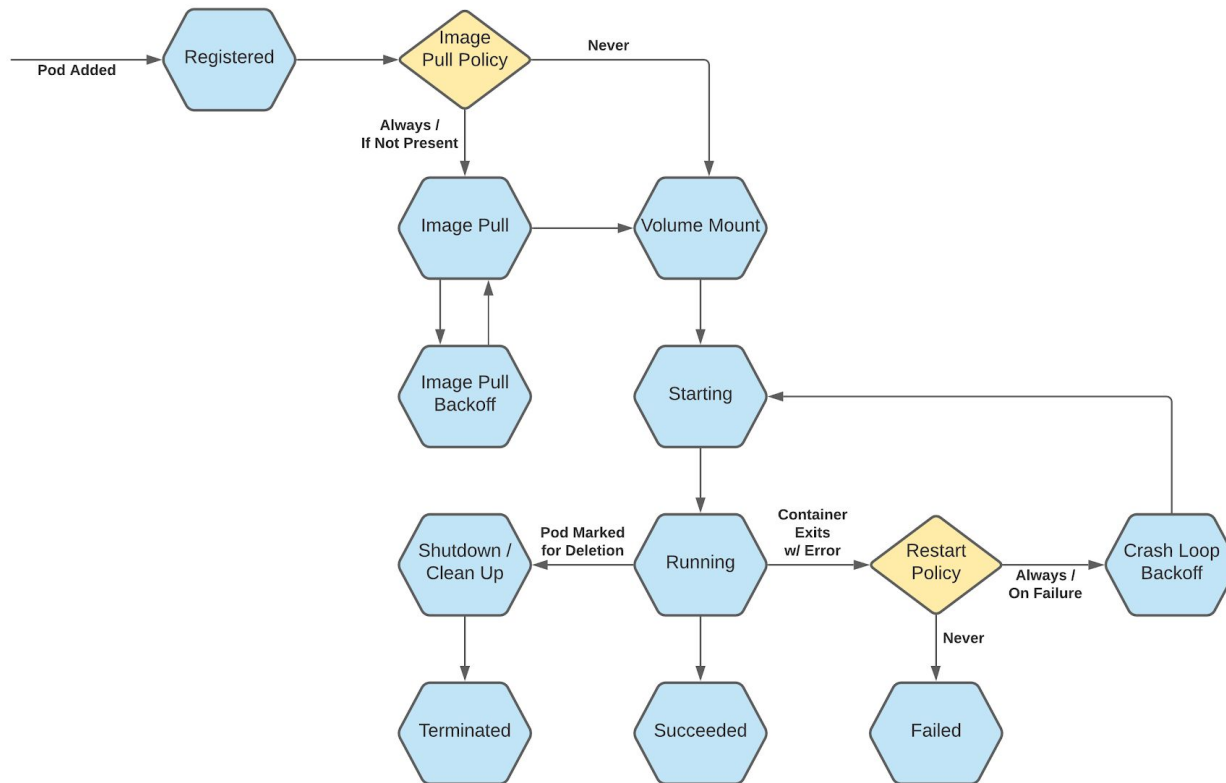
KubeCon



CloudNativeCon

North America 2020

Virtual



Rust State Machine



KubeCon



CloudNativeCon

North America 2020

Virtual

- [A Fistful of States: More State Machine Patterns in Rust](#)
- Flexible framework for implementing Kubelet control loop.
- Enforced at compile time:
 - Valid states
 - Valid state transitions
- Automatic Pod status updates
- Error handling within context of control loop.

Krustlet Architecture



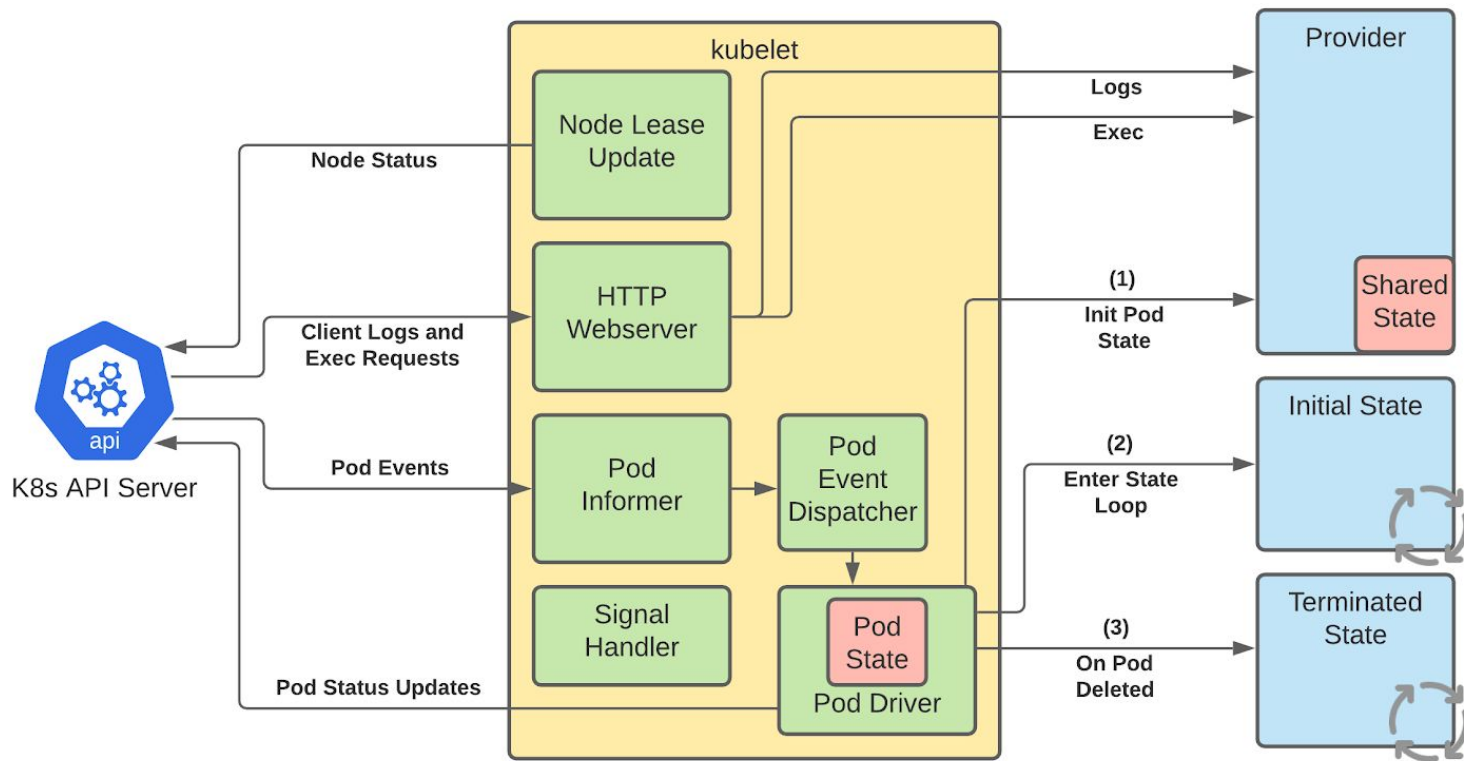
KubeCon



CloudNativeCon

North America 2020

Virtual



Conclusion



KubeCon



CloudNativeCon

North America 2020

Virtual

- Kubelet communication patterns
- Pod behavior
- Rust
- Shout outs
 - Taylor Thomas
 - Matt Fisher
 - Ivan Towlson
- Contributing to Krustlet

