



Writing Service Mesh Controllers in Rust

Eliza Weisman, Linkerd maintainer

Eliza Weisman

Software Engineer, Buoyant

 @mycoliza

 @hawkw

 @eliza





Rust is **fast** and **memory safe**, all **without garbage collection**



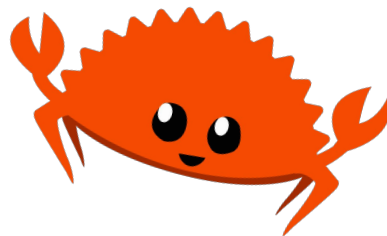
Ultralight, ultrafast, security-first **service mesh** for Kubernetes.

- ★ Created by **Buoyant**
- ★ **Goals:** secure, efficient, fast
- ★ **Data plane:** proxies application traffic
- ★ **Control plane:** tells the proxies what to do



Linkerd + Rust

- **Pure Rust data plane** since the release of Linkerd 2
- First **Rust control plane component** released in Linkerd 2.11



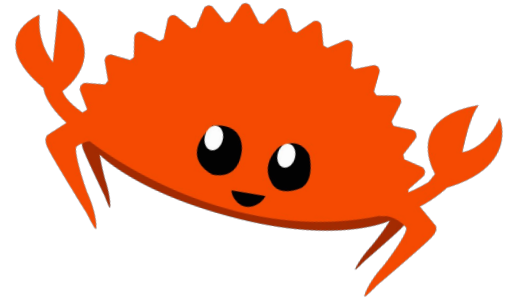
Rust was **the only choice** for the **proxy**.

We **didn't want GC pauses**...

...but we **didn't want C++ CVEs**, either.

Rust isn't the only choice for the **control plane**, but **it is a good one** (and we like it).

Rust has great **language features**
dev tools
libraries



Linkerd's control plane **watches**
Kubernetes resources and **serves gRPC**
APIs for the proxies

The **policy controller** indexes **policy resources** and **associates them with ports** on pods.

Writing **controllers** in Rust means **talking to Kubernetes**, which means **bindings for the Kubernetes API**.

crates.io [Browse All Crates](#) | Eliza Weisman ▾

kube v0.75.0

Kubernetes client and async controller runtime

[#kubernetes](#) [#runtime](#) [#client](#)

[Readme](#) 92 Versions Dependencies Dependents

kube-rs

crates.io v0.75.0 MSRV 1.60 MK8SV v1 20 opensf best practices passing
chat 1697 online

A [Rust](#) client for [Kubernetes](#) in the style of a more generic [client-go](#), a runtime abstraction inspired by [controller-runtime](#), and a derive macro for [CRDs](#) inspired by [kubebuilder](#). Hosted by [CNCF](#) as a [Sandbox Project](#)

These crates build upon Kubernetes [apimachinery](#) + [api concepts](#) to enable generic abstractions. These abstractions allow Rust reinterpretations of reflectors, controllers, and custom resource interfaces, so that you can write applications

Metadata

📅 6 days ago
📄 Apache-2.0
📦 10.1 kB

Install

Add the following line to your Cargo.toml file:

```
kube = "0.75.0"
```

Documentation

kube + rt = kubert



Docs for `kubert` 0.11.0

Platform Feature flags Releases Rust Find crate

```
pub use self::server::ServerArgs;
```

Modules

<code>admin</code>	Admin server utilities.
<code>client</code>	Utilities for configuring a <code>kube_client::Client</code> from the command line
<code>errors</code>	Utilities for handling errors
<code>index</code>	Utilities for maintaining a shared index derived from Kubernetes resources.
<code>initialized</code>	A utility for waiting for components to be initialized.
<code>log</code>	Configures the global default tracing subscriber
<code>queue</code>	A bounded, delayed, multi-producer, single-consumer queue for deferring work in response to scheduler updates.
<code>runtime</code>	A controller runtime
<code>server</code>	Helpers for configuring and running an HTTPS server, especially for admission controllers and API extensions
<code>shutdown</code>	Drives graceful shutdown when the process receives a signal.

Structs

<code>LogInitError</code>	Error returned by <code>try_init</code> if a global default subscriber could not be initialized.
---------------------------	--



```
pub trait IndexNamespacedResource<T> {  
    fn apply(&mut self, resource: T);  
    fn delete(&mut self, namespace: String, name: String);  
    fn reset(  
        &mut self,  
        resources: Vec<T>,  
        removed: HashMap<String, HashSet<String>>,  
    );  
}
```



```
apiVersion: policy.linkerd.io/v1beta1
kind: Server
metadata:
  namespace: linkerd-viz
  name: admin
spec:
  podSelector:
    matchLabels:
      linkerd.io/extension: viz
  port: admin-http
  proxyProtocol: HTTP/1
```

```
#[derive(Clone, Debug, PartialEq, Eq)]
#[derive(kube::CustomResource)]
#[derive(Deserialize, Serialize, JsonSchema)]
#[kube(
  group = "policy.linkerd.io",
  version = "v1beta1",
  kind = "Server",
  namespaced
)]
#[serde(rename_all = "camelCase")]
pub struct ServerSpec {
  pub pod_selector: labels::Selector,
  pub port: Port,
  pub proxy_protocol: Option<ProxyProtocol>,
}
```

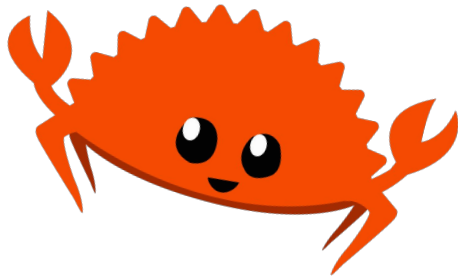
```
#[tokio::main]
async fn main() -> Result<(), Box<dyn Error>> {
    // Build the Kubert runtime (Kubert also has arg-parsing helpers)
    let Args { client } = Args::parse();
    let runtime = kubert::Runtime::builder()
        .with_client(client)
        .build()
        .await?;

    // Start indexing the Server CRD
    let index = Index::default();
    let servers = runtime.watch_all::<Server>(ListParams::default());
    tokio::spawn(kubert::index::namespaced(index.clone(), servers));

    // Do more stuff, like spawning gRPC servers...

    runtime.run().await.map_err(Into::into)
}
```


<https://github.com/linkerd/linkerd2/tree/main/policy-controller>



BUOYANT'S

SERVICEMESH
ACADEMY

The creators of
LINKERD

Monthly hands-on, engineer-focused training from the creators of the service mesh

Nov 11-17: What really happens at startup: a deep dive into Linkerd, init containers, CNI plugins, and more



SIGN UP TODAY!

buoyant.io/sma

Fully managed LINKERD on any Kubernetes cluster

Buoyant Cloud automated upgrades, data plane version tracking, mesh health alerts, and much, much more.

BOOK A DEMO

buoyant.io/demo



Thanks so much!

Eliza Weisman

Software Engineer, Buoyant

eliza@buoyant.io

 [@BuoyantIO](https://twitter.com/BuoyantIO)

 buoyant.io