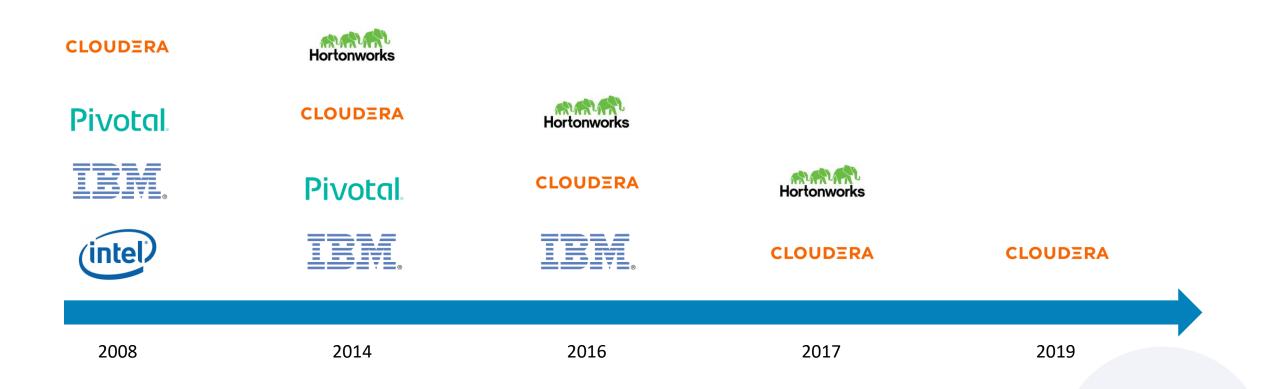


Building a New Big Data Distribution Based on Kubernetes – With a Twist!

15th June 2021 | Track Scale | Format Talk Intermediate

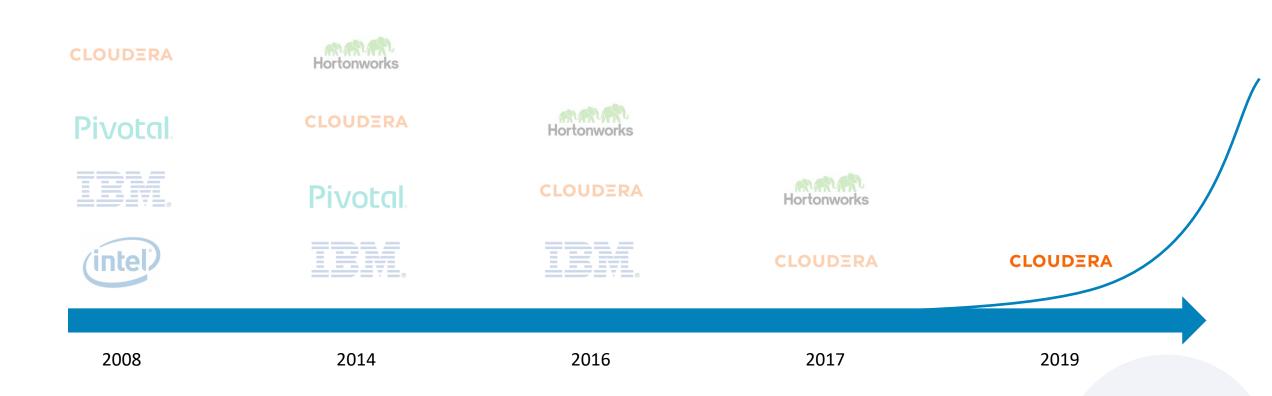


Hadoop Distributions over Time



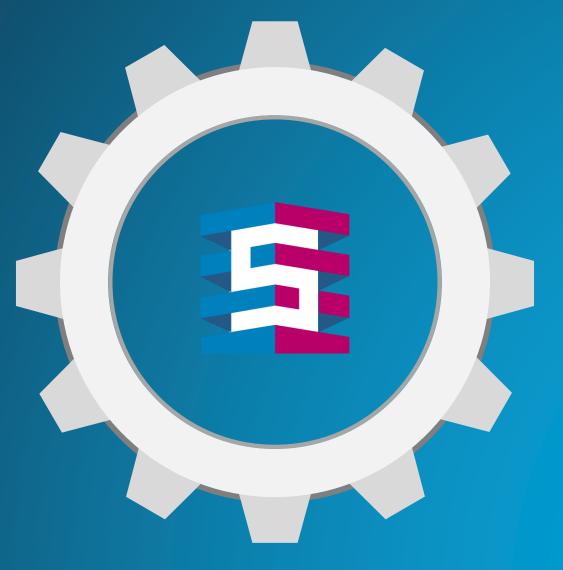


How Did This Impact the per Node Price?





The History of Stackable





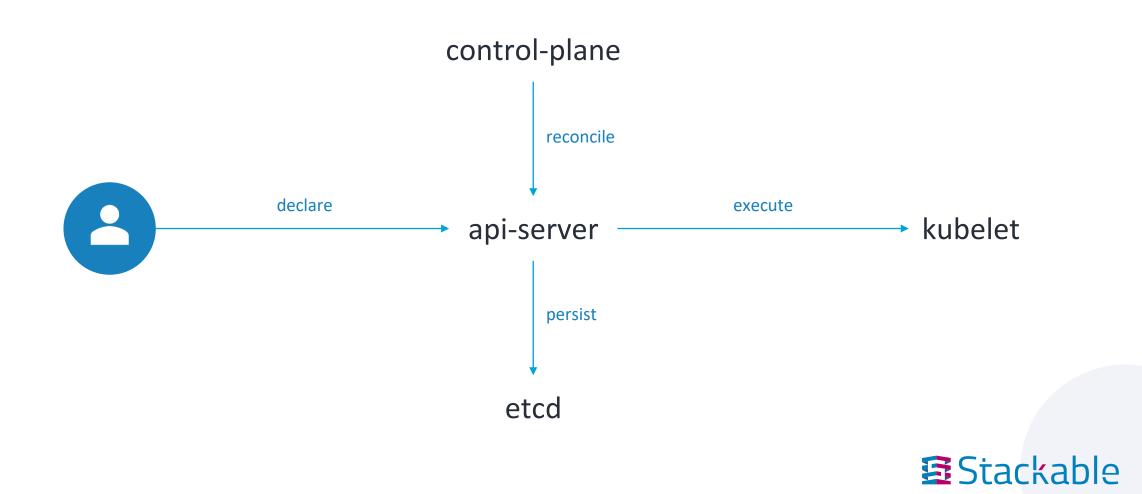
Kubernetes on One (Well, Two) Slide(s)

An orchestration framework to manage services Declarative by nature you define what reality should look like and Kubernetes goes and makes it happen (This has some shortcomings, more on that later)

Kubernetes has built-in data types Extensible via Custom Resource Definitions



Kubernetes from Outer Space





Stackable & Kubernetes



Stackable & Kubernetes



We wanted to build a distribution on top of Kubernetes.

Customers told us "no".



 \gg

So, we came up with a way to use the best of both worlds.

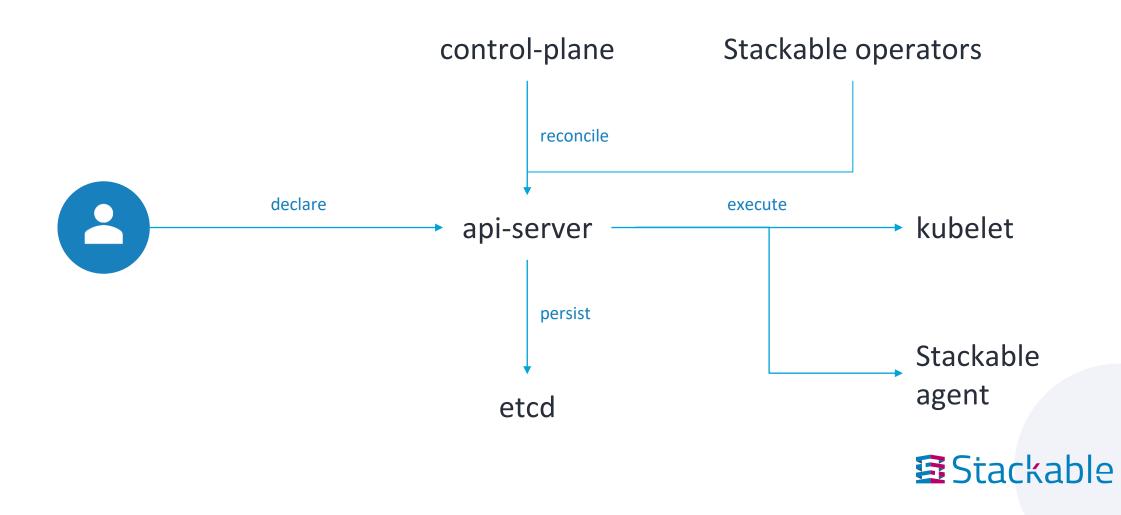
Use the Kubernetes control plane and concepts (e.g. Operators,GitOps, ...) but allow reusing existing admin knowledge (e.g. systemd, OS Packages, ...)

This allows hybrid scenarios:

- Part of the stack on "bare metal"
- Part in native Kubernetes in containers



Stackable & Kubernetes



Stackable & Kubernetes

We only use a few core concepts at the moment:

- Pod
- Node
- ConfigMap
- Secret
- ...add more as needed

kcp is a minimal Kubernetes API server

O Build passing

How minimal exactly? kcp doesn't know about Pod S or Node S, let alone Deployment S, Service S, LoadBalancer S, etc.

By default, kcp only knows about:

- Namespace S
- ServiceAccount s and role-based access control types like Role and RoleBinding
- Secret s and ConfigMap s, to store configuration data
- CustomResourceDefinition s, to define new types
- a handful of other low-level resources like Lease S, Event S, etc.

• • •				
\$ kubectl get pods				
error: the server doesn't ha	ave a resource			
<pre>\$ kubectl api-resources</pre>				
NAME	SHORTNAMES	APIVERSION	NAMESPACED	
configmaps			true	ConfigMap
events				
limitranges			true	LimitRange
namespaces				Namespace
resourcequotas				
			true	
				ServiceAccount
customresourcedefinitions	crd, crds	apiextensions.k8s.io/v1		CustomResourceDefinition
apiservices		apiregistration.k8s.io/v1		APIService
tokenreviews		authentication.k8s.io/v1		TokenReview
localsubjectaccessreviews		authorization.k8s.io/v1	true	LocalSubjectAccessReview
selfsubjectaccessreviews		authorization.k8s.io/v1		SelfSubjectAccessReview
		authorization.k8s.io/v1		SelfSubjectRulesReview
subjectaccessreviews				SubjectAccessReview
certificatesigningrequests		certificates.k8s.io/vlbetal		CertificateSigningRequest
		events.k8s.io/v1beta1	true	Event
clusterrolebindings				ClusterRoleBind
		rbac.authorization.k8s.io/v1		
rolebindings				RoleBinding
roles		rbac.authorization.k8s.io/v1	true	Role

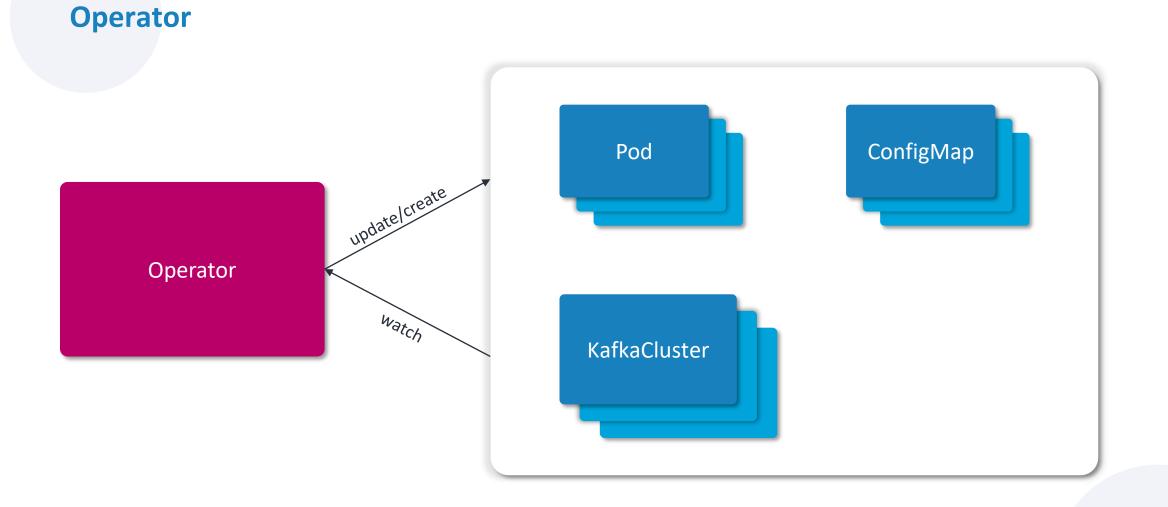




Operator

```
let operator = extract_brain(human);
operator.run(my_cluster);
```







Operator Best Practices

- Reconcile loop
- Common labels
- Common nomenclature
- Common status & events
- Common configuration/CRDs
- Common monitoring, tracing, metrics
- Hacks: Imperative commands



Our distribution should feel consistent hence we decided to write our own framework.

Stackable 5

Reconcile Loops – Two Alternative Styles



1





Stackable Stackable

Stackable Operator Framework



Stackable Operator Framework

Lots of convenience functions

- Retrieve existing Pods
- Retrieve/Set conditions
- Remove pods
- • •

Often used Higher Level Abstractions

- Remove unneeded pods
- Find nodes which should have pods

This requires a few conventions about what our objects look like!



So, How Does It Work in Real Life?

self.init status()

- .then(self.context.delete illegal pods(ContinuationStrategy::AllRequeue))
 - .then(self.context.wait_for_terminating_pods())
 - .then(self.context.wait for running and ready pods())
 - .then(self.context.delete excess pods(ContinuationStrategy::OneRequeue))
 - .then(self.create missing pods())



Why Rust?

Most Kubernetes code & 3rd-party operators are written in Go. Why did you choose Rust?

- 🕅 We didn't know either language
- 🖑 We tried both
- We looked at the library ecosystem
- \bigcirc We liked Rust. End of story.
- \wp If you want details:
 - Error Handling
 - Enums
 - Generics
 - No GC
 - Security
- Go's structural typing is sweet



Rust Kubernetes Ecosystem - k8s-openapi

This crate is a Rust Kubernetes API client. It contains bindings for the resources and operations in the Kubernetes client API, auto-generated from the OpenAPI spec.

This crate is *not* generated using Swagger or the OpenAPI Generator directly, as clients generated by the common client generator are. This gives this crate a few important advantages.

More Details on https://github.com/Arnavion/k8s-openapi





Rust Kubernetes Ecosystem - kube-rs



and a derive macro for CRDs inspired by kubebuilder.

More Details on https://github.com/clux/kube-rs





Rust Kubernetes Ecosystem - Krustlet

Crustlet: Kubernetes Kubelet in Rust for running WASM

Trustlet 1.0 coming soon!

Krustlet acts as a Kubelet by listening on the event stream for new pods that the scheduler assigns to it based on specific Kubernetes tolerations.

The default implementation of Krustlet listens for the architecture wasm32-wasi and schedules those workloads to run in a wasmtime -based runtime instead of a container runtime.

More Details on https://github.com/deislabs/krustlet









Lars Francke ars.francke@stackable.de +49 4103 926 3100 www.stackable.de

Contact me on

>BLN BZZ/ WRDS

Thank You!

Stackable



Sönke Liebau soenke.liebau@stackable.de +49 4103 926 3100 www.stackable.de

Contact me on

♡ ¥ in X