

Charliecloud as a Kubernetes Container Runtime

Angelica Loshak

Rensselaer Polytechnic
Institute

angelicaloshak@gmail.com

London Bielicke

University of California, Los Angeles
Rhodes College

bielickelondon@gmail.com

Mentors

Reid Priedhorsky

Lucas Caudill



Purpose

HPC users seek new systems to support the increasing demand for novel workflows, including those for AI applications

How can we efficiently manage large-scale, distributed and containerized environments?



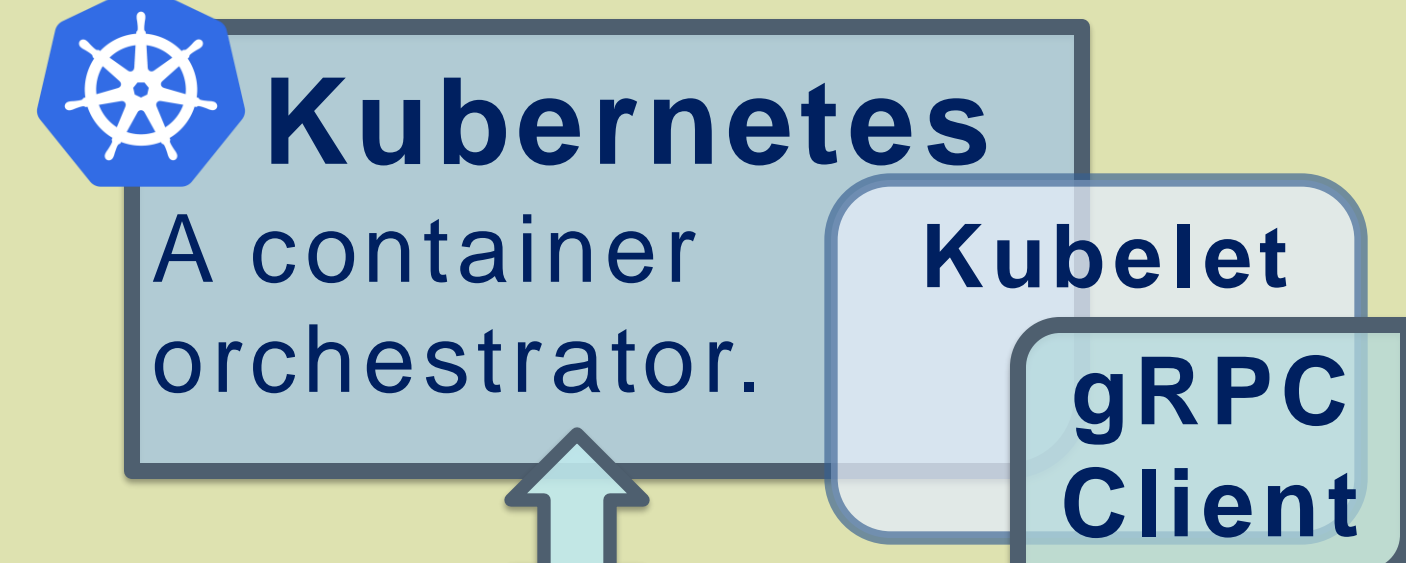
kubernetes

“Charliecloud is a lightweight, fully unprivileged container runtime.”

-Reid Priedhorsky

CRI (Container Runtime Interface)

Enables interaction between Kubernetes and the container runtime.



Users specify a container runtime such as Charliecloud.

We implemented the CRI as a gRPC Remote Procedure Calls (gRPC) server.

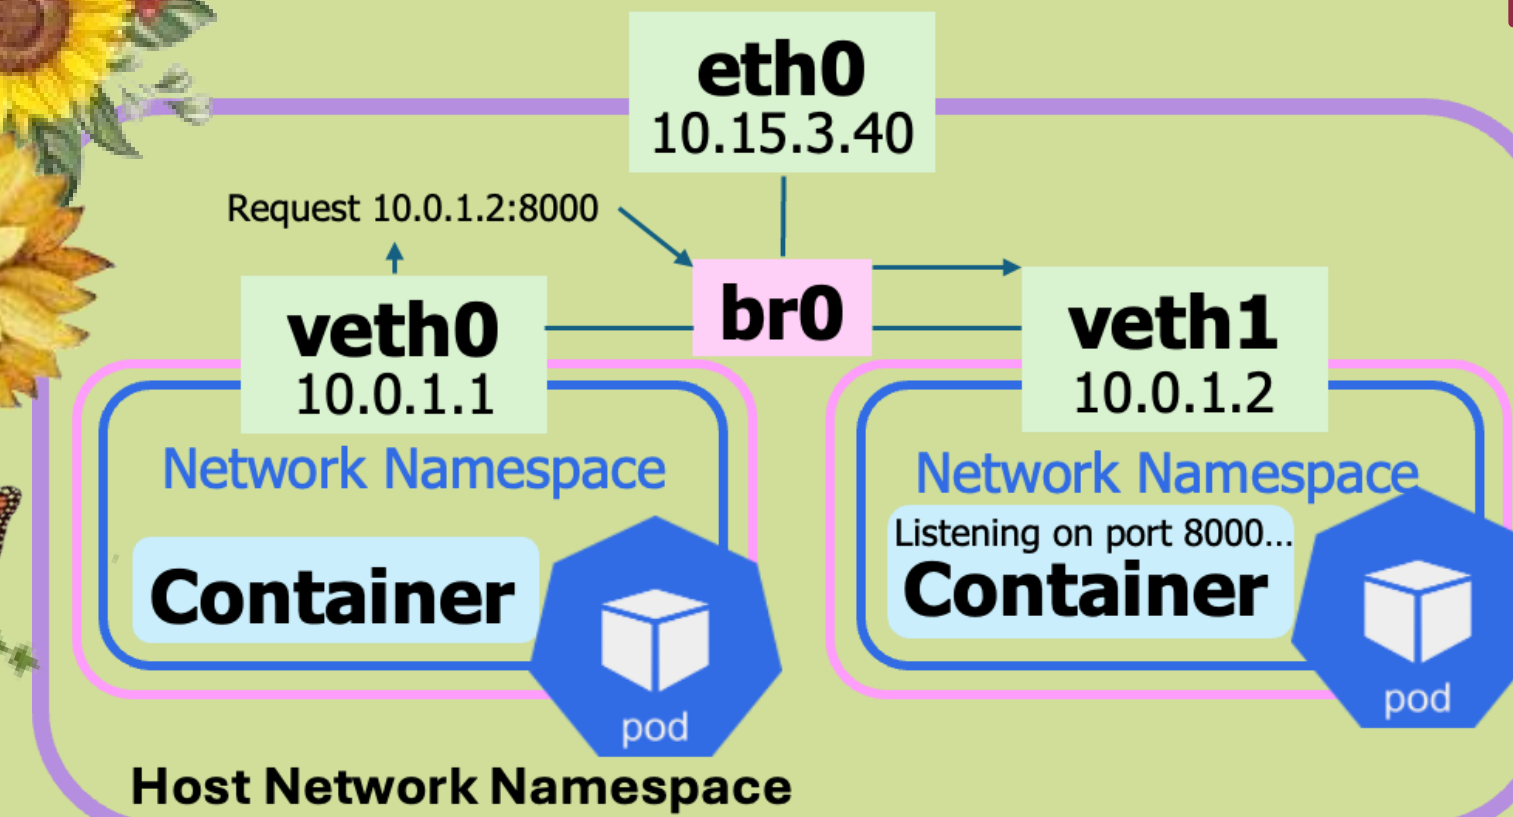
The Kubelet runs on each node and makes calls to the container runtime over a Unix socket.

gRPC: Google's Remote Procedure Call framework

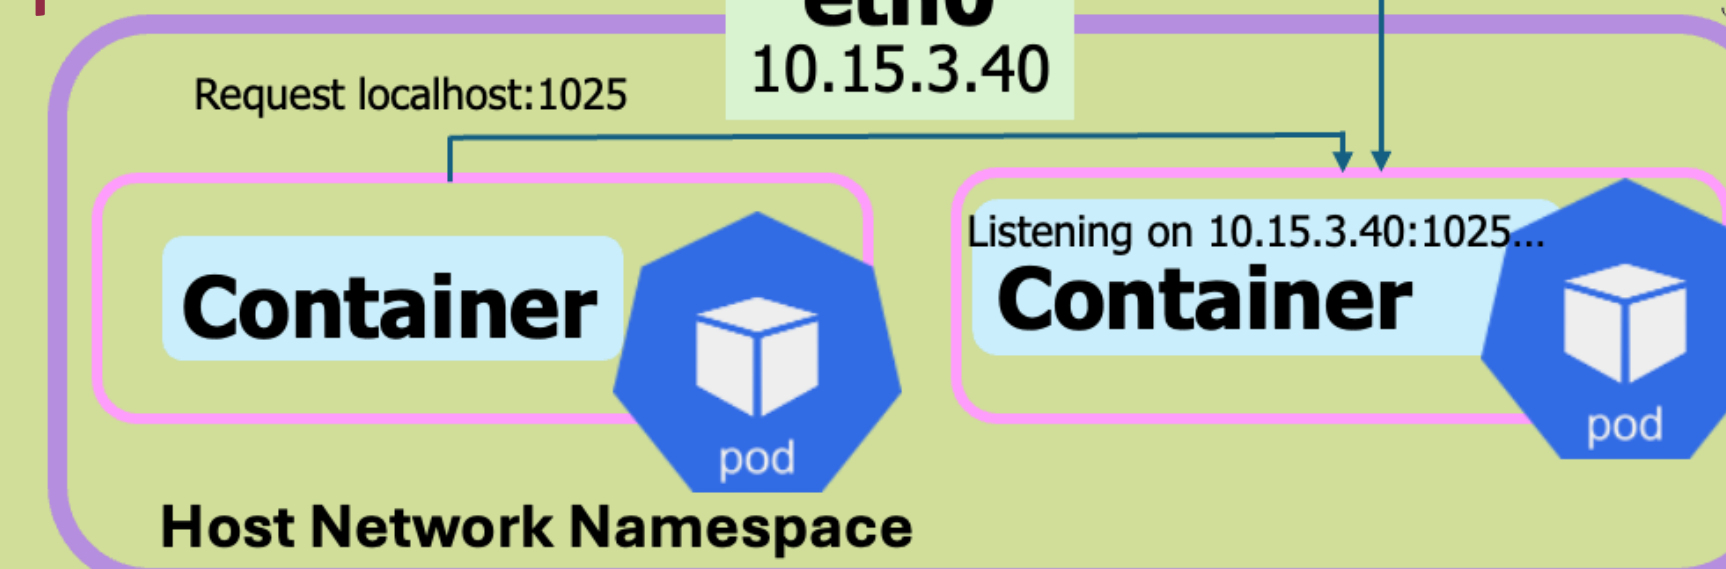
Here the gRPC server is listening on the UNIX socket...

Server listening on socket unix:///tmp/al.sock...

CRI Assumptions



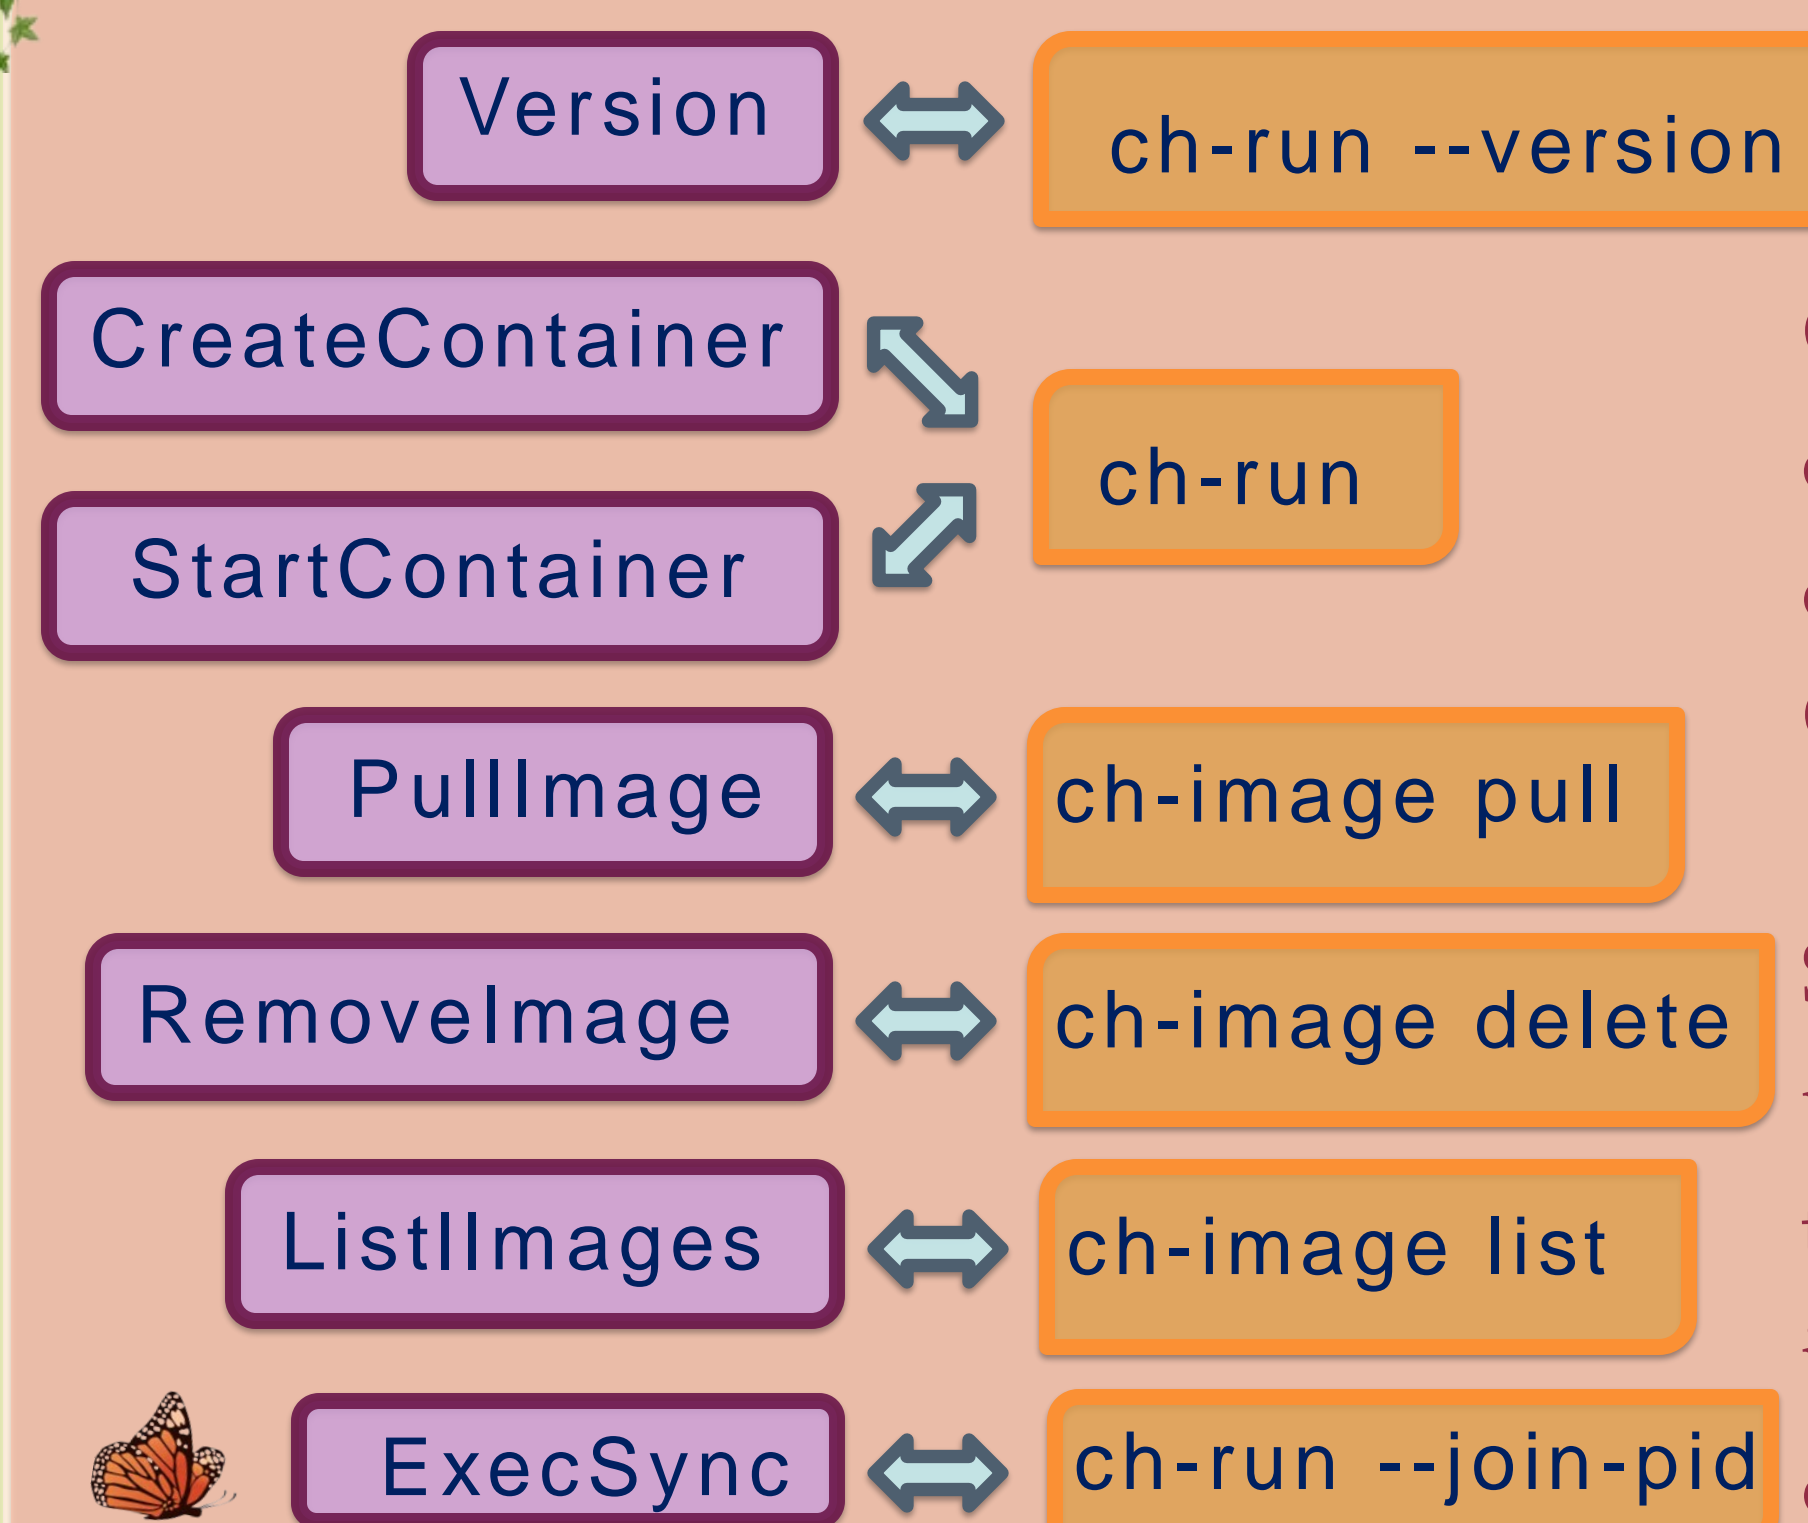
Users can change Kubernetes config to use `hostNetwork: true` and modify the port numbers.



Charliecloud containers:

- 1) Share the host's Linux network namespace
- 2) Use unprivileged ports.

Charliecloud Compatibility



Charliecloud stores container states during CreateContainer().

StartContainer() uses ch-run under the hood, using the image and CMD defined in state.

Testing

25/94 CRI tests pass,

more methods still need implementation:

```
25 Passed | 54 Failed | 0 Pending | 15 Skipped
TestCRISuite (881.83s)
```

Command executed in the Kubernetes container:

```
ch-run[2509125]: executing: sh -c "echo hello World; sleep 1000"
```

Charliecloud + Kubernetes can:

1. Create and start containers.
2. Track container health metrics.
3. Run pods.
4. Respawn containers in a pod.

Events:	Type	Reason	Age	From	Message
	Normal	Pulled	4m39s (x12 over 72m)	kubelet	Container image "ra
	Normal	Created	4m39s (x12 over 72m)	kubelet	Created container m
	Normal	Started	4m39s (x12 over 72m)	kubelet	Started container m

Table 1: shows the kubelet pulling an image, creating a container, starting, and respawning a container using Charliecloud

Kubelet: communicates between control plane and Kubernetes nodes

We demonstrated that Kubernetes and Charliecloud are compatible tools.

Charliecloud hosts a gRPC server with 700 lines of code. Less than 50 lines of modification to other Charliecloud source code.

The Kubernetes containers are running:

CONTAINER	IMAGE	STATE
c2d5645d-b73e	rancher/local-path-provisioner:v0.0.	Running
9ce425fc-9cf6	rancher/mirrored-coredns-coredns:1.1	Running
cfedf840-4528	rancher/mirrored-metrics-server:v0.7	Running

Future is unreachable...or is it?

We need to finish the implementations of: HostNetwork

to communicate between Pods on different nodes

Container configuration options:

runwithgid, Readiness probes, etc.

Attach and Exec methods:

Build on Exec and attach to stream over http.

