



CRI-O: All the runtime k8s needs, and nothing more



CRI

- In the beginning, there was docker.
- Then along came rkt.
- Adding and maintaining new runtimes in k8s was becoming a burden.
- Hence CRI was introduced to decouple k8s from container runtimes.



Why CRI-0

- Was it possible to create a minimal runtime using just pluggable standard components?



CRI-O Scope

- Scope is tied to the **CRI**
- **Shaped** around Kubernetes
- Only supported user is Kubernetes



Overview of components

- **OCI compatible container runtimes** are supported. Tested with runc and Clear Containers.
- **github.com/containers/storage** is used for managing layers and creating root filesystem
 - Overlay, devicemapper, aufs, btrfs (defaulting to overlay).
 - NFS support on the way.
- **github.com/containers/image** is used for pulling images from registries
 - Battle-tested: we embed it in our docker fork and that's used by OpenShift for pulling images through docker. We use it for image signature verification as well.
 - Supports docker schema2 version 1 and version 2 - covers all corner cases and passes docker and kubernetes tests.

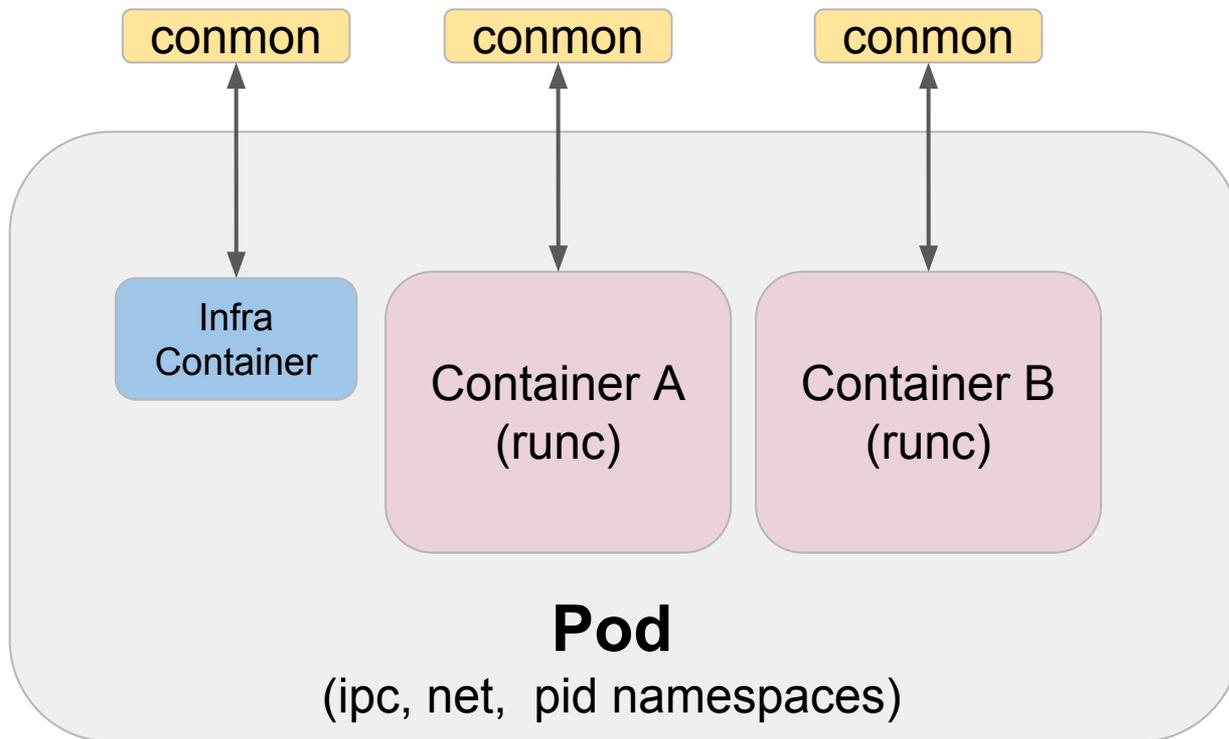


Overview of components (continued)

- **oci-runtime-tools** library is used to generate OCI configs for containers
- **CNI** is used for setting up networking
 - Tested with Flannel, Weave and openshift-sdn
- **common** is a utility for:
 - Monitoring
 - Logging
 - Handling tty
 - Serving attach clients
 - Detecting and reporting OOM

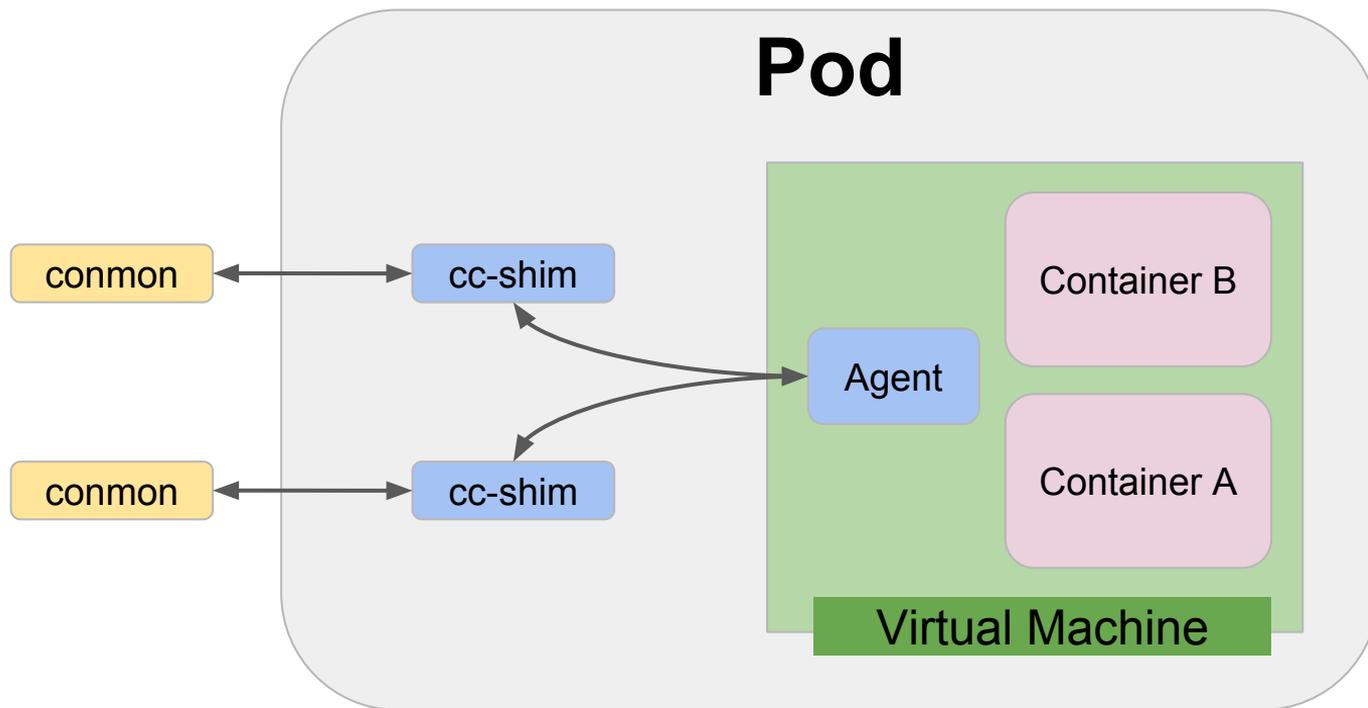


Pod architecture (runc)



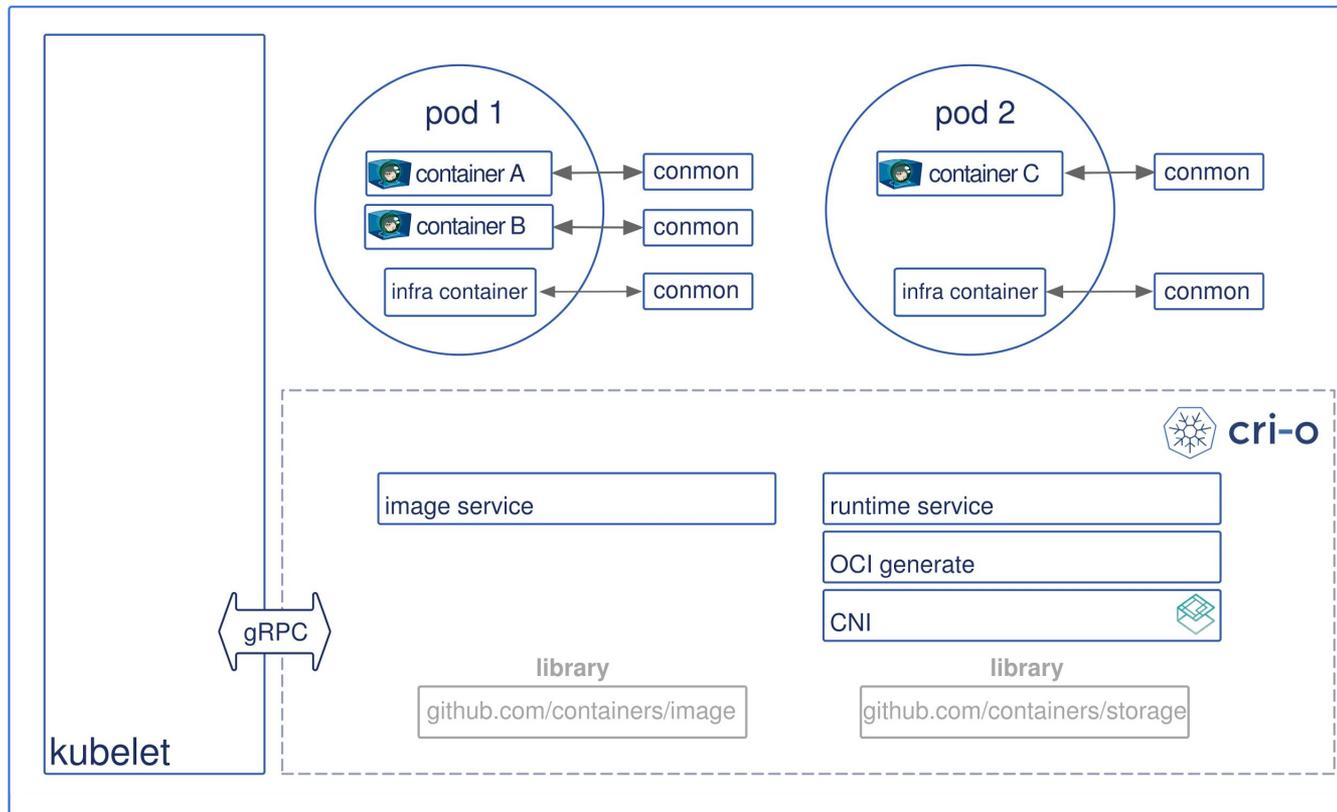


Pod architecture (Clear Containers)





Architecture





Status

- **All** node conformance tests passing. (These are run on each PR)
- **All** e2e tests passing.
- **All** critests passing.
- **All** CRI APIs implemented.
- **All** kubernetesbyexample.com examples work
- Master is tracking kube 1.9
- Maintainers/contributors from **Red Hat, Intel & SUSE**
- **Kubeadm** works for setting up k8s with CRI-O (supports Ubuntu/Centos/Fedora)
- **Minikube** integration is in progress
- Support for **mixed workloads** (runc/Clear Containers)



CRI-O versioning

- 1.0 -> kube 1.7
- 1.8 -> kube 1.8
- 1.9 (master) -> kube 1.9



Demo

- Pods w/ logs
- Jobs w/ logs
- Exec
- Attach/Detach
- Openshift Routes
- Kubeadm and mixed workloads - <https://asciinema.org/a/123891>
- Multi node/Multi OS cluster - <https://asciinema.org/a/124131>



kpod-Management tool for containers and images

Daemonless tool to debug cri-o

Based on Docker CLI



Next steps

- Releasing 1.9
- Graduating out of incubator
- Tracking and supporting k8s versions



Contribute/Further Info

Blog: <https://medium.com/cri-o>

Github: <https://github.com/kubernetes-incubator/cri-o>

IRC: freenode: #cri-o

Site: <https://cri-o.io>

Questions?