



CRI-O: OCI-based Kubernetes Runtime

Mrunal Patel & Dan Walsh (Red Hat)

OpenShift Commons Briefing
August 2017



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.

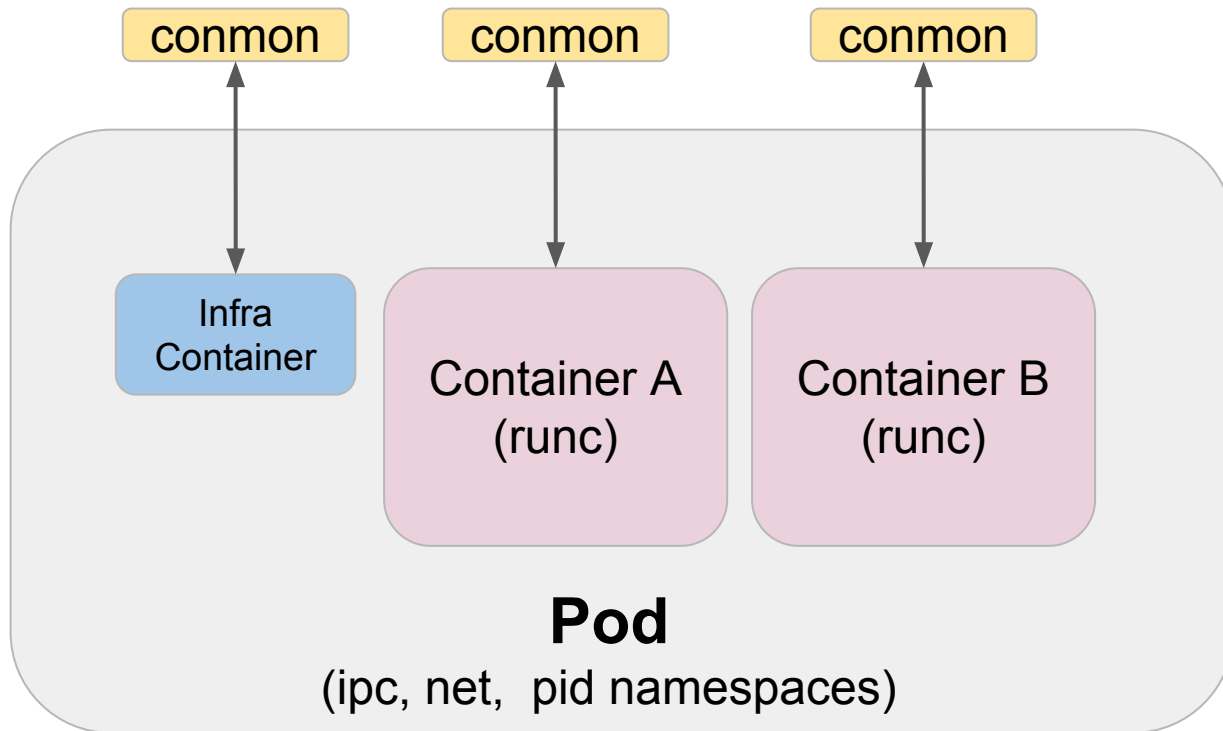


Overview of components (continued)

- **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

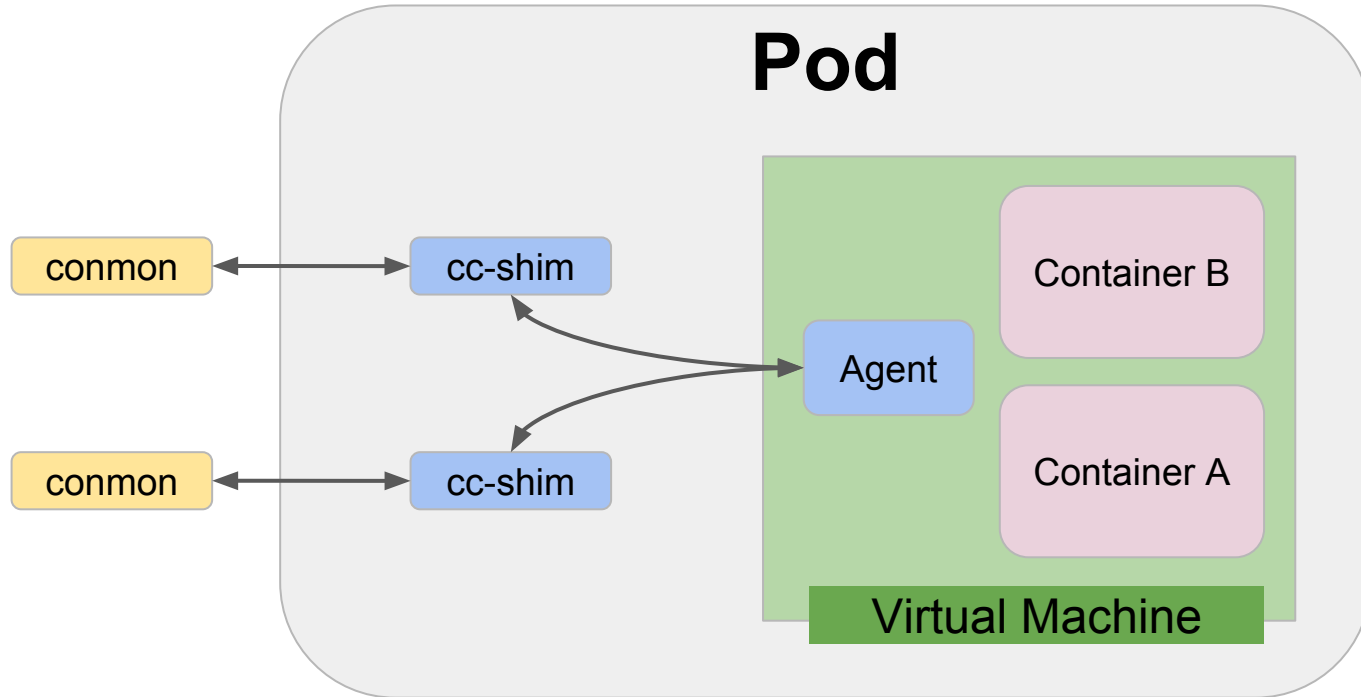


Pod architecture (runc)



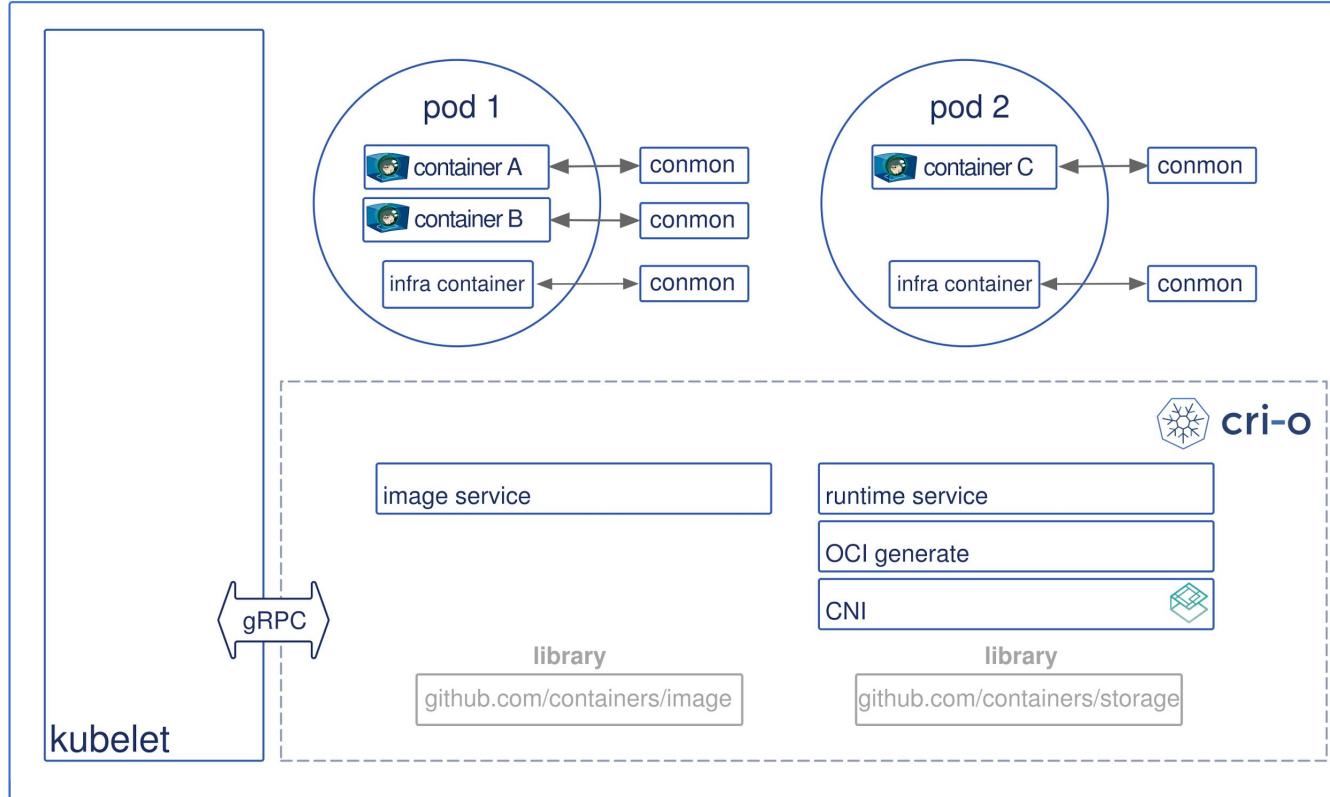


Pod architecture (Clear Containers)





Architecture





Status

- **All** node conformance tests passing. (These are run on each PR)
- **All** e2e tests passing.
- **All** CRI APIs implemented.
- **All** kubernetesbyexample.com examples work
- 1.0.0.beta was released last week
- 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)



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

Commands

Command	Description
crio(8)	OCI Kubernetes Container Runtime daemon
kpod(1)	Simple management tool for pods and images
kpod-cp(1)	Copy files/folders between a container and the local filesystem
kpod-diff(1)	Inspect changes on a container or image's filesystem
kpod-export(1)	Export container's filesystem contents as a tar archive
kpod-history(1)	Shows the history of an image
kpod-images(1)	List images in local storage
kpod-info(1)	Display system information
kpod-inspect(1)	Display the configuration of a container or image
kpod-load(1)	Load an image from docker archive or oci
kpod-mount(1)	Mount a working container's root filesystem
kpod-pull(1)	Pull an image from a registry
kpod-push(1)	Push an image to a specified destination
kpod-rmi(1)	Removes one or more images
kpod-save(1)	Saves an image to an archive
kpod-tag(1)	Add an additional name to a local image
kpod-umount(1)	Unmount a working container's root filesystem
kpod-version(1)	Display the Kpod version information



Next steps

- Releasing 1.0.0
- Graduating out of incubator
- Tracking and supporting k8s versions
- Target OpenShift 3.7 and OpenShift Online



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?