


rkt Introduction



What is rkt?

Iago López Galeiras, Kinvolk
Alban Créquy, Kinvolk



About us


- Living in Berlin
- Kinvolk co-founders
- Worked a lot on rkt and rktlet
- Currently maintainers of rkt and rktlet

```
* github: https://github.com/iaguis  
          https://github.com/alban  
* twitter: @iaguis  
          @albcr
```

Plan


- What is rkt?
- How does rkt work? (rkt internals)
- What's new?
- What's missing?





a modern, secure, composable container runtime





an implementation of appc (image format,
execution environment)



a standard application container
open specification
associated tooling




- github.com/appc/spec ("appc spec")
- github.com/appc/acserver
- github.com/appc/docker2aci
- [github.com/~~appc~~/containernetworking/cni](https://github.com/appc/containernetworking/cni)
- github.com/appc/...

appc spec in a nutshell

- Image Format (ACI)
 - what does an application consist of?
- Image Discovery
 - how can an image be located?
- Pods
 - how can applications be grouped and run?
- Executor (runtime)
 - what does the execution environment look like?





simple CLI tool
golang + Linux
self-contained
init system/distro agnostic

simple CLI tool

no (mandatory) daemon

apps run directly under spawning process



simple CLI tool

```
bash/systemd/runit
├─ rkt (stage0)
│   └─ pod (stage1)
│       ├── app1 (stage2)
│       └── app2 (stage2)
```

- stage0
 - rkt binary (CLI)
 - discover, fetch, manage application images
 - set up pod filesystems
- stage1
 - exec environment for pods
 - process lifecycle management
 - resource constraints
- stage2
 - actual application



stage1 (swappable)

- default implementation
 - Linux namespaces + cgroups for isolation
 - based on `systemd-nspawn` + `systemd`
- kvm implementation
 - hardware virtualization for isolation
 - based on `qemu` + `systemd`
- fly
 - no isolation
 - just a `chroot`
- others



rkt: recent developments

- Share host IPC namespace
- Fixes to make rktlet possible
 - Bugfixes, especially within the rkt experiments
 - Integration between RKT_EXPERIMENT_ATTACH and k8s
- Switched default kvm flavor from lkvm to qemu
- Bugfixes



rkt: recent developments (2)

- rkt documentation

Filters ▾ **Labels** **Milestones**

Clear current search query, filters, and sorts

0 Open 31 Closed Author ▾ Labels ▾



rkt: what's missing?

- Update CNI version to 0.6.0
- Stabilize rkt experiments and remove RKT_EXPERIMENT_* env variables
- Experiment with using runc to set up stage2 runtime
- Possible native OCI support
- General bugfixing :)



Demo

rkctl k8s cluster on kube-spawn

Thank You!

Questions?