



Container Live Migration

Adrian Reber

Tübix 2019
July 06

Definition:

Container Live Migration

Transfer Running Container

Serialize on Source System

Transfer to Destination System

Checkpoint/Restore in Userspace

CRIU

Multiple Integrations Exist

CRIU

First Step: Checkpointing

Collect Details From

```
/proc/<PID>/*
```

Seize Process Using `ptrace()`

Parasite Code

Parasite Code

Most favorite part

Parasite Code

And the craziest

Parasite Code

Injected into the process

Parasite Code

Daemon waiting for commands

Parasite Code

Removed after usage

Checkpointing Finished

Checkpointing Finished

All relevant information written

Checkpointing Finished

Target process is killed

Checkpointing Finished

Or continues to run

Container Live Migration

SELinux

Second/Last Step: Restoring

Read Checkpoint Images

clone() For Each PID/TID

CRIU Morphs Itself

Open and position file descriptors

CRIU Morphs Itself

Map memory pages

CRIU Morphs Itself

Load security settings

CRIU Morphs Itself

Jump into restored process

Container Live Migration

Container Live Migration

OpenVZ

Container Live Migration

LXC/LXD

Container Live Migration

Docker

Container Live Migration

Podman

Podman: daemonless

Podman: rootless

Podman: Checkpoint/Restore

October 2018

Podman: Checkpoint/Restore

Required runc and CRIU changes

Podman: Container Live Migration

June 2019

Podman: Container Live Migration

Required runc, CRIU, SELinux
changes


```
1 # podman run --rm -d adrianreber/wildfly-hello
2 699f33eb7fecbc5bbb00400be0aa79c888dbc63a54cac7bd2eed836a57d8a68a
3 # podman inspect -l --format "{{.NetworkSettings.IPAddress}}"
4 10.88.0.247
5 # curl 10.88.0.247:8080/helloworld/
6 0
7 # curl 10.88.0.247:8080/helloworld/
8 1
9 # podman container checkpoint -l --export=/tmp/chkpt.tar.gz
10 699f33eb7fecbc5bbb00400be0aa79c888dbc63a54cac7bd2eed836a57d8a68a
11 # scp /tmp/chkpt.tar.gz rhel08:/tmp
```

```
1 # podman container restore --import=/tmp/chkpt.tar.gz
2 699f33eb7fecbc5bbb00400be0aa79c888dbc63a54cac7bd2eed836a57d8a68a
3 # podman inspect -l --format "{{.NetworkSettings.IPAddress}}"
4 10.88.0.247
5 # curl 10.88.0.247:8080/helloworld/
6 2
7 # curl 10.88.0.247:8080/helloworld/
8 3
```

```
1 # podman container restore --import=/tmp/chkpt.tar.gz -n hello1
2 d02feec894d77f66cc82484fe77ae369396a85f6d05594dc156c21e685942dd
3 # podman container restore --import=/tmp/chkpt.tar.gz -n hello2
4 735efb4fee6961d3eee069beb28dde5cbc6fc46c1a32a43ecc993d04c02015b2
5 # podman inspect --format "{{.NetworkSettings.IPAddress}}" hello1
6 10.88.0.248
7 # podman inspect --format "{{.NetworkSettings.IPAddress}}" hello2
8 10.88.0.249
9 # curl 10.88.0.248:8080/helloworld/
10 2
11 # curl 10.88.0.249:8080/helloworld/
12 2
```

<https://lisas.de/~adrian/container-live-migration-article.pdf>

<https://asciinema.org/a/249922>

<https://asciinema.org/a/249918>

<https://lisas.de/~adrian/posts/2019-Apr-10-criu-and-selinux.html>

<https://criu.org/Podman>

https://twitter.com/adrian__reber

