



Container Migration with Kubernetes and CRIU

Adrian Reber
Principal Software Engineer

2021, May 14
<https://people.redhat.com/~areber/2021-lpd.pdf>

Agenda

Background

Use cases

Details

Demos

Future

Checkpoint/Restore in Userspace

CRIU

Definition:

Container Live Migration

Transfer Running Container

Serialize on Source System

Transfer to Destination System

Background

Kubernetes

Container

Container namespaces

Container cgroups

Container seccomp

Container Runtime

runc Or crun

Container Engine


CRI-O or containerd or Podman


Container Orchestration

Kubernetes


Use Cases

Reboot and Save State








Container





Quick Startup











Container Live Migration







CRIU

First Step: Checkpointing

Seize Process Using `ptrace()`

Collect Details From

`/proc/<PID>/*`

Parasite Code

Parasite Code

Most favorite part

Parasite Code

And the craziest

Parasite Code

Injected into the process

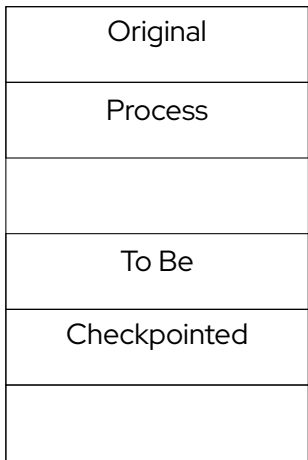
Parasite Code

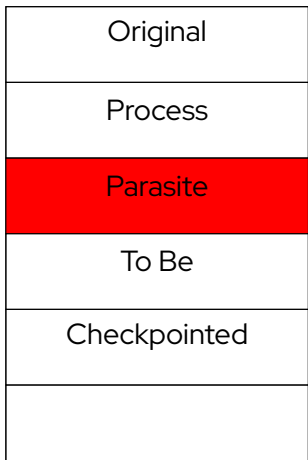
Daemon waiting for commands

Parasite Code

Removed after usage

Original
Process
Code
To Be
Checkpointed





Original
Process
Code
To Be
Checkpointed

Checkpointing Finished

Checkpointing Finished

All relevant information written

Checkpointing Finished

Target process is killed

Checkpointing Finished

Or continues to run

Second/Last Step: Restoring

Read Checkpoint Images

`clone()` For Each PID/TID

`clone3()` with Linux 5.5

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

```
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 d02feeec894d77f66cc82484fe77ae369396a85f6d05594dc156c21e685942dd
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
```

Future:

```
kubect1 migrate
```

Future:

Non-root checkpoint/restore

`CAP_CHECKPOINT_RESTORE`

Summary

- CRIU can checkpoint and restore containers
- Integrated in different containers engines
- Used in production
- Reboot into new kernel without losing container state
- Start multiple copies
- Migrate running containers



Thank you