CRIU - Checkpoint Restore in Userspace

Fabio Andrijauskas –
fandrijauskas@ucsd.edu


ESC
:q!

COPYING LINE OF CODE IN:

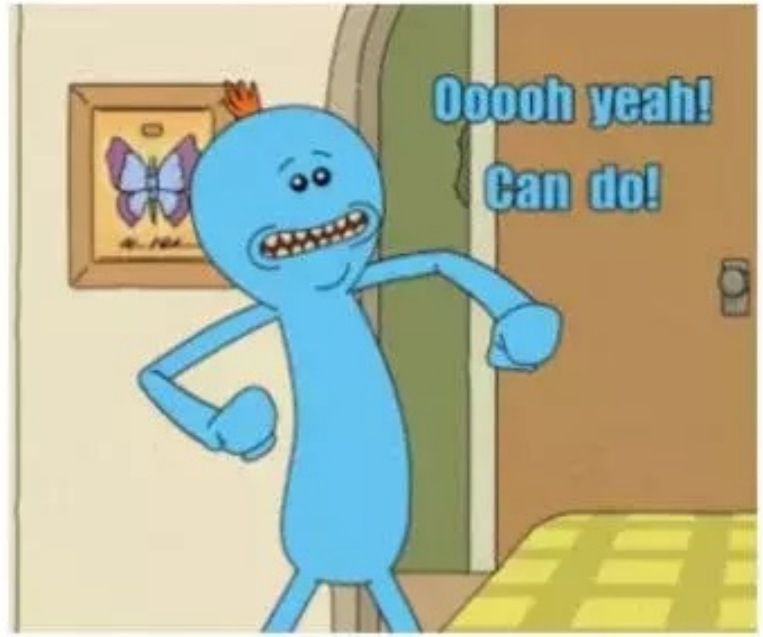

COPYING LINE OF CODE IN:


What's the best code editor?
VSCode
Sublime
VIM!!
imgflip.com

Checkpoint

- While executing scientific applications, they must sometimes stop the process due to hardware problems or even job end-of-life.

- The checkpoint can be provided by the application or by an external tool.

- It is not an easy task to checkpoint software using an external tool.

CRIU

- While executing scientific applications, they must sometimes stop the process due to hardware problems or even job end-of-life.
- CRIU (Checkpoint Restore in Userspace, pronounced kree-oo) is a tool for checkpointing and restoring applications in GNU/Linux environment.
- https://criu.org

CRIU, OSG, and more

- The object is to check the features of CRIU:
  - Check the basic features.
  - How will CRIU behave on a batch system?
  - Check the networks and **containers** support.

Tests

- Simple C code: writing on the terminal. Using the same operational system and hardware
  - The test was successful, only requiring a few options on restore and dump.

- LAMMPS: read and write files. Using the same operational system and hardware
  - If CRIU detects any file or directory change location, the restoration fails. To restore, it is required to set the paths from before the CRIU dump.

Tests

- Simple C code: dump from Intel hardware to AMD hardware.
- Simple C code – dump from AMD hardware to Intel hardware.
  - If any processer code optimization is used, it will fail, and all the system libraries should be the same.
- Open files and dump the software
  - The file structure should be the same on dump and restore

Tests

- **TCP and UDP Network connections:**
  - **One host: client and server stay on the same host.**
  - **Two hosts – need to stay on the same host.**

Tests

- **Nvidia GPUs are not supported at all.**
- **Singularity, Podman, and docker – Test to dump and restore inside the container**
- **Singularity – Test to dump and restore outside the container**
- **Podman – Test to dump and restore outside the container**
  - Using podman/CRIU integration interface is possible to dump and check containers using the podman command.

Tests

- **Dumping and restoring a POD using CVMFS**
- **Dumping and restoring a process with FORK**
- **Dumping and restoring a process with MPI**

Conclusions

- CRIU can provide several options to stop and restore applications.
- It is possible to control applications with multiple threads and processes maintaining network connections.
- CRIU supports containers using docker and podman.

Conclusions

- It requires a “form” of root access: sudo, SUID Bit, or Kernel capabilities.
- To use Kernel capabilities requires a specific version of CRIU and Linux.
- Restoring a previously checkpointed process requires the same directory paths used during restoration as during checkpointing.

https://path-cc.io/GIL/criu_checkpoint_restore_userspace/

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.