



Contribution ID: 249

Type: **not specified**

News from academia: FatELF, RDMA and CRIU

Tuesday, November 13, 2018 9:15 AM (30 minutes)

As a part of ongoing research project we have added several features to CRIU: post-copy memory migration, post-copy migration over RDMA and support from cross-architecture checkpoint-restore.

The “plain” post-copy migration is already upstream and, up to hiccups that regularly show up in CI, it can be considered working so there is not much to discuss about it.

The post-copy migration over RDMA aims to reduce the remote page fault latency. We have a working prototype that replaces TCP transport for memory transfer with RDMA. We still do not have solid performance evaluation, but if RDMA will provide the expected reduction in page fault latency, we are going to work with the CRIU community to make the RDMA support upstream.

The cross-architecture checkpoint-restore is the most peculiar and controversial feature. Various aspects of heterogeneous ISA execution have been a hot research topic for a while, and we decided to see what it would take to make CRIU capable of migrating an application between architectures.

The idea is simple: if we create the binary for all architectures so that all the symbols have exactly the same address, then restoring a dump created on a different architecture is a matter of transforming the stack and the registers.

This transformation relies on the metadata generated by the specialized compiler that generates multiple object files from a single source (one for each architecture), hence the FatELF.

Up till now we were able to force CRIU to create a checkpoint of an extended “Hello, World” application on arm64 and restore this application on x86.

I agree to abide by the anti-harassment policy

Presenters: NIDER, Joel (IBM); RAPOPORT, Mike (IBM)

Session Classification: Containers MC