



Contribution ID: 194

Type: **not specified**

Preserving guest memory across kexec

Monday, September 12, 2022 12:15 PM (25 minutes)

Live update is a mechanism to support deploying updates to a running hypervisor in a way that has limited impact to virtual machines. This is done by pausing the virtual machines, stashing KVM state, kexecing into a new kernel, and restarting the VMM process. The challenge is guest memory: how can it be preserved and restored across kexec?

This talk describes a solution to this problem: moving guest memory out of the kernel managed domain, and providing control of memory mappings to userspace. Userspace is then able to restore the memory mappings of the processes and virtual machines via a FUSE-like interface for page table management.

We describe some requirements, options, why the FUSE-style options was chosen, an an overview of the work-in-progress implementation. Opinions are collected around other use cases this functionality could support. Next steps around finalising the design and working to get this included upstream are discussed.

This is a follow-on the the initial RFC presented at LSF-MM a few months ago: <https://lwn.net/SubscriberLink/895453/71c46dbe09426f59/>

I agree to abide by the anti-harassment policy

Yes

Primary author: GOWANS, James (Amazon EC2)

Co-author: WOODHOUSE, David

Presenter: GOWANS, James (Amazon EC2)

Session Classification: Kernel Memory Management MC

Track Classification: LPC Microconference: Kernel Memory Management MC