



Contribution ID: 363

Type: **not specified**

Transparent Huge-Page support for guest_memfd

Supporting hugepages is critical to unlocking the performance potential of confidential VMs. While there has been a lot of recent development activity around adding hugepage support to guest_memfd to support HugeTLB-like use-cases, there are also efforts underway to enable hugepage support for guest_memfd via the normal buddy allocator and leveraging Transparent Huge-Pages to provide hugepage support for cases where reserving pools of memory for HugeTLB pages can be prohibitive.

One notable aspect of the guest_memfd design for hugepages is that it will rely heavily on the ability to split hugepages on-demand for accounting purposes, which brings about many similarities, but also some notable differences, with how splitting/merging of THP pages is currently handled by the kernel outside of guest_memfd. For instance, guest_memfd would benefit heavily from being able to merge THPs without freeing them to the kernel, and this may have additional use-cases outside of guest_memfd. This session will focus on highlighting cases like these and exploring where it makes sense to unify handling vs. special-casing things for guest_memfd.

Primary author: ROTH, Michael (AMD)

Presenter: ROTH, Michael (AMD)

Session Classification: Kernel Memory Management MC

Track Classification: Kernel Memory Management MC