



Contribution ID: 133

Type: **not specified**

Pageable-based Virtual Machine (PVM) as a PV flavor for KVM

Thursday 19 September 2024 10:00 (30 minutes)

Nowadays, there are various needs to run a VM in the public cloud, such as running a security container to isolate workloads or encapsulating an application into a VM for migration or rapid kernel testing utilizing cost-effective spot VMs. However, nested virtualization on KVM requires hardware support and is usually disabled by the cloud provider for safety reasons. Additionally, the current nested architecture involves complex and expensive transitions between the L0 hypervisor and L1 hypervisor. Therefore, we are introducing a new virtualization framework built upon the KVM hypervisor that does not require hardware-assisted virtualization techniques. This framework serves as a PV flavor for KVM, allowing the running of a VM in the public cloud without nested virtualization support.

We have provided our RFC patch set to present the PVM design. In this session, we plan to share various use cases of PVM and present its inherent value. In particular, we want to discuss the underlying technology associated with the x86 subsystem and the KVM subsystem, and the aspects of these subsystems improved in the implementation of PVM (e.g. improved shadow paging which also helps for nested TDP), as well as the collaboration to consolidate the work for the future. Moreover, we would also like to address the security model (e.g. side channel attacks between guest/host) in PVM compared to other PV virtualization implementations. Furthermore, we aim to explore the possibility of extending PVM to more architectures (e.g. ARM64 and RISC-V) to establish it as a common PV flavor for KVM.

Primary authors: JIANGSHAN, Lai (AntGroup); HOU, Wenlong

Presenters: JIANGSHAN, Lai (AntGroup); HOU, Wenlong

Session Classification: KVM Microconference

Track Classification: KVM Microconference