

Linux Plumbers Conference 2024



Contribution ID: 33

Type: **not specified**

KVM Microconference

CFP closes on July 12th.

KVM (Kernel-based Virtual Machine) enables the use of hardware features to improve the efficiency, performance, and security of virtual machines (VMs) created and managed by userspace. KVM was originally developed to accelerate VMs running a traditional kernel and operating system, in a world where the host kernel and userspace are part of the VM's trusted computing base (TCB).

KVM has long since expanded to cover a wide (and growing) array of use cases, e.g. sandboxing untrusted workloads, depriving third party code, reducing the TCB of security sensitive workloads, etc. The expectations placed on KVM have also matured accordingly, e.g. functionality that once was "good enough" no longer meets the needs and demands of KVM users.

The KVM Microconference will focus on how to evolve KVM and adjacent subsystems in order to satisfy new and upcoming requirements. Of particular interest is extending and enhancing guest_memfd, a guest-first memory API that was heavily discussed at the 2023 KVM Microconference, and merged in v6.8.

The KVM MC is expected to have strong representation from maintainers (KVM and non-KVM), hardware vendors (Intel, AMD, ARM, RISC-V, etc), cloud (AWS, Google, Oracle, etc), client (Android, ChromeOS), and open source stalwarts such as Red Hat and SUSE.

Potential Topics:

- Removing guest memory from the host kernel's direct map[1]
- Mapping guest_memfd into host userspace[2]
- Hugepage support for guest_memfd[3]
- Eliminating "struct page" for guest_memfd
- Passthrough/mediated PMU virtualization[4]
- Pagetable-based Virtual Machine (PVM)[5]
- Optimizing/hardening KVM usage of GUP[6][7]
- Live migration support for guest_memfd
- Defining KVM requirements for hardware vendors
- Utilizing "fault" injection to increase test coverage of edge cases

[1] <https://lore.kernel.org/all/cc1bb8e9bc3e1ab637700a4d3defec95b55060a.camel@amazon.com>

[2] <https://lore.kernel.org/all/20240222161047.402609-1-tabba@google.com>

[3] https://lore.kernel.org/all/CABgObfa=DH7FySBviF63OS9sVog_wt-AqYgtUAGKqnY5Bizivw@mail.gmail.com

[4] <https://lore.kernel.org/all/20240126085444.324918-1-xiong.y.zhang@linux.intel.com>

[5] <https://lore.kernel.org/all/20240226143630.33643-1-jiangshanlai@gmail.com>

[6] <https://lore.kernel.org/all/CABgObfZCay5-zaZd9mCYGMeS106L055CxsdOWWvRTUk2TPYycg@mail.gmail.com>

[7] <https://lore.kernel.org/all/20240320005024.3216282-1-seanjc@google.com>

Primary authors: BONZINI, Paolo (Red Hat, Inc.); BONZINI, Paolo (Red Hat); CHRISTOPHERSON, Sean (Google)

Presenters: BONZINI, Paolo (Red Hat, Inc.); BONZINI, Paolo (Red Hat); CHRISTOPHERSON, Sean (Google)

Track Classification: LPC Microconference Proposals