

# Multi-KVM

# Introduction

Multi-KVM is a proposal to allow multiple, independent KVM modules to be loaded, unloaded, and run concurrently on the same Linux host to:

- Upgrade and rollback KVM without disrupting running VMs.
- Enable running KVM modules with different parameters on the same host.
- Facilitate easier A/B testing for KVM.

# Design Proposal

## Isolation

- Hide KVM internals from kernel at-large
- <https://lore.kernel.org/lkml/20230916003118.2540661-1-seanjc@google.com>

## Multi-KVM

- Collapse `kvm_intel.ko` and `kvm_amd.ko` into `kvm.ko`
- Expose multiple `kvmN.ko` modules and `/dev/kvmN` devices
- “N” is defined via a Kconfig string (null/off by default)

## VAC

- Extract shared system resources out of KVM into a new “base module” (VAC)
- VAC = Virtualization Acceleration Code (Unupgradable Units Module)

## User Space

- No user space *VMM* changes required, e.g. symlinks (`/dev/kvm` → `/dev/kvmN`), bind mounts, etc.
- User space ultimately controls deployment, usage, lifecycles, etc.

# Key Feedback

- Hiding KVM internals from the rest of the kernel
- Base module name / namespace
- Collapsing x86 vendor modules into kvm.ko
- ???