



Contribution ID: 166

Type: **not specified**

Multikernel Architecture: Kernel-to-Kernel Isolation and Resource Management

Thursday 11 December 2025 10:00 (45 minutes)

Modern Linux faces fundamental scaling challenges with shared resource contention, noisy neighbor effects, and monolithic kernel constraints. VMs provide isolation but impose significant hypervisor overhead, while containers share kernel vulnerabilities and lack performance isolation.

We propose the multikernel architecture enabling multiple isolated Linux kernel instances on a single machine. A privileged host kernel dynamically spawns independent kernel instances, each with dedicated CPU cores, memory regions, and I/O hardware resources. Our implementation extends existing kernel mechanisms: dynamic kernel spawning using enhanced kexec for on-demand instantiation without system reboot, hardware resource partitioning through fine-grained CPU/memory/device isolation, inter-kernel communication via IPI and shared memory regions, and live resource migration enabling runtime resource reassignment for zero-downtime upgrades.

Each spawned kernel runs standard Linux userspace, delivering strong isolation at near-native performance while maintaining the complete Linux compatibility. More importantly, comparing with KVM/virtio stack, multikernel is significantly simpler thus has great potential.

Multikernel architecture enables seamless application deployment with complete workload isolation, specialized kernel optimization, and enhanced security boundaries suitable for cloud multi-tenancy and safety-critical systems.

Primary author: WANG, Cong

Presenter: WANG, Cong

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track