

Linux Plumbers Conference 2024



Contribution ID: 399

Type: **not specified**

Fast, Flexible, and Practical Kernel Extensions

Friday 20 September 2024 10:30 (30 minutes)

The ability to safely extend OS kernel functionality is a longstanding goal in OS design, with the widespread use of the eBPF framework in Linux and Windows only underlining the benefits of such extensibility. However, existing approaches to kernel extensibility constrain users in the extent of functionality that can be offloaded to the kernel or the performance overheads incurred by their extensions.

We present KFlex: an approach that provides an improved tradeoff between the expressibility and performance of safe kernel extensions. KFlex separates the enforcement of kernel safety from the enforcement of extension correctness, and uses bespoke mechanisms for each to enable users to express diverse functionality in their extensions at low runtime overheads. We implement KFlex in the context of the Linux kernel, and our prototype is fully backward compatible with existing eBPF-based extensions, making it immediately useful to practitioners. Our evaluation demonstrates that KFlex enables offloading functionality that cannot be offloaded today, and in doing so, provides significant end-to-end performance benefits for latency-sensitive applications.

Primary author: DWIVEDI, Kumar Kartikeya (EPFL)

Co-authors: IYER, Rishabh (UC Berkeley); Prof. KASHYAP, Sanidhya (EPFL)

Presenter: DWIVEDI, Kumar Kartikeya (EPFL)

Session Classification: eBPF Track

Track Classification: eBPF Track