



Contribution ID: 400

Type: **not specified**

eBPF-specialized Kernel for I/O Intensive Applications

Friday 20 September 2024 16:00 (30 minutes)

Datacenter workloads have demanding performance requirements, including the need for high throughput and low tail latency while maintaining high server utilization. While modern hardware is compatible with these goals, overheads and inefficiencies in today's operating systems remain a critical bottleneck. Several research proposals aim to address this problem by designing dataplane OSES with specialized I/O stacks and scheduling algorithms. However, these proposals have poor backwards compatibility, and lack broader hardware support. There also exist proposals which use eBPF to offload application logic into the kernel, and bypass the OS layers to provide better performance. However, these lose multi-tenancy and isolation between competing workloads, and require rewriting applications.

In this talk, we will discuss an approach that provides similar performance, isolation, and multiplexing benefits as existing proposals while retaining the backwards compatibility and reliability of the Linux I/O stack. Our approach is based on the observation that the bottlenecks in the Linux I/O stack are due to how execution resources are allocated to each stage of I/O processing, and not due to inefficiencies in individual I/O stack components. So, we believe eBPF-driven extensibility can address this by allowing workload-specific specialization of the I/O data path, balancing the functional and performance requirements datacenter workloads. Finally, eBPF programs can provide the necessary abstractions to stitch together request processing logic of applications across user and kernel mode, providing an end-to-end solution.

Primary author: DWIVEDI, Kumar Kartikeya (EPFL)

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

Presenter: DWIVEDI, Kumar Kartikeya (EPFL)

Session Classification: eBPF Track

Track Classification: eBPF Track