



Contribution ID: 344

Type: **not specified**

cache_ext: Customizing the Page Cache with eBPF

The page cache is central to the performance of many applications. However, its one-size-fits-all eviction policy may perform poorly for many workloads. While the systems community has experimented with new and adaptive eviction policies in non-kernel settings (e.g., key-value stores, CDNs), it is very difficult to implement such policies in the kernel. We design a flexible eBPF-based framework for the Linux page cache, called `cache_ext` that allows developers to customize the page cache without modifying the kernel.

We demonstrate the flexibility of `cache_ext`'s interface by using it to implement eight different policies, including sophisticated eviction algorithms. Our evaluation finds that it is indeed beneficial for applications to customize the page cache to match their workloads' unique properties, and that they can achieve up to 70% higher throughput and 58% lower tail latency.

Primary authors: ZARKADAS, Ioannis (Columbia University); ZUSSMAN, Tal (Columbia University)

Co-authors: CHENG, Andrew (Columbia University); CIDON, Asaf (Columbia University); FRANKE, Hubertus (IBM Research); CARIN, Jeremy (Columbia University); PFEFFERLE, Jonas (IBM Research)

Presenter: ZUSSMAN, Tal (Columbia University)

Session Classification: eBPF Track

Track Classification: eBPF Track