



Contribution ID: 194

Type: **not specified**

Fingerprinting GPU workloads with eBPF

Bridging the Observability Gap: Using eBPF for GPU Workload Identification

Modern computing workloads are increasingly offloaded to GPUs, yet our ability to observe and understand the specific tasks running on these accelerators from the host kernel remains limited. This fundamental lack of visibility hinders system administrators, security engineers, and resource schedulers. While existing solutions often rely on application-level telemetry or proprietary vendor tools, they fail to provide a holistic, kernel-level view of GPU activity.

This paper introduces a novel solution that leverages eBPF to gain unprecedented insight into GPU workloads. By attaching eBPF tracepoints to the NVIDIA kernel driver, we can non-intrusively monitor and profile the sequence of driver function calls for any given task. This method captures a rich set of metrics—including the frequency and timing of calls—to build a unique “execution fingerprint” for the workload.

We demonstrate that these profiles are distinct and reproducible. Using a variety of real-world use cases, including machine learning training, inference, and cryptocurrency mining, we show that our eBPF-based approach can reliably identify the underlying workload with high accuracy.

Our findings highlight the power of eBPF as a versatile and potent tool for bridging the critical observability gap between the host kernel and accelerated compute devices.

Primary author: GOGELA, Jiri (Trend Micro)

Presenter: GOGELA, Jiri (Trend Micro)

Session Classification: eBPF Track

Track Classification: eBPF Track