Monitoring Clusters Using Extended Berkeley Packet Filter (eBPF)

Abstract

Matthew Aiden Phillips, Alexis Malee Ng, Weston James Cadena
Mentors: Travis Cotton, Nick Jones, and Jonathan Nielsen

LA-UR-23-28964

Modern High Performance Computing (HPC) systems vary immensely in size, software, and purpose; creating vast differences in production workflows and performance. Troubleshooting system calls, both in kernel and user space, within HPC has traditionally not been straightforward. With metrics that were previously impossible to view, eBPF delivers highly detailed and objective information clarifying system performance [3].

eBPF tools are used for performance analysis and observability at a low cost and overhead by attaching probes to kernel and user space system calls. These probes provide non-invasive access into kernel and user routines without disrupting processes. eBPF Compiler Collection (BCC) [4] is a set of tools derived from eBPF that provide insight into the kernel and user stack and software applications using tracing, sampling, and snooping. These tools could be used for improved system monitoring, application profiling, and general troubleshooting. This provides users avenues to improve simulation efficiency and give administrators more in-depth answers about HPC systems.

This project is aimed to determine the viability of BCC tools in an HPC environment. With BCC tools, we were able to characterize the following:

- NFS Latency and Bandwidth
- CPU stacks and nature of running processes
- Intercommunication of parallel processes over TCP
- Cache hits and misses over time

Los Alamos
NATIONAL LABORATORY