Forensic performance analysis

Forensic performance analysis

Keywords : performance analysis, cloud, system programming, multithreading

Context

When running an application in production at large scale (typically a cloud-based web service), you cannot afford to collect too much performance data because of the overhead of data collection. Thus, when a performance problem happens (eg. a client request takes too long to process), investigating is complex because of the lack of information. Reproducing the bug may also be unfeasible as it may requires many production servers running a particular workload and hitting a race condition.

Goals

The goal of this project is to design a tracing tool tailored for **forensic performance analysis**. This tracing tool would run on production servers at large scale without degrading performance (tolerance: 1% overhead), and without generating large amounts of data. As the application runs, the tracing tool collects performance data in circular buffers. When a performance problem happens, the collected data is dumped into a file for post-mortem investigation.

The main steps of this project are:

- designing low-overhead mechanism for collecting performance data in circular buffers
- evaluate the overhead of the mechanism, and how many events per second can be collected without degrading performance
- implement a data dump mechanism able to extract data while the application runs
- evaluate the overhead of the data dump mechanism

Contact

François Trahay francois.trahay@telecom-sudparis.eu
Parallel & Distributed Systems group
Télécom SudParis