

Multi-Granularity Memory Mirroring via Binary Translation in Cloud Environments

ABSTRACT:

As the size of DRAM memory grows in clusters, memory errors are common. Current memory availability strategies mostly focus on memory backup and error recovery. Hardware solutions like mirror memory needs costly peripheral equipments while existing software approaches reduce the expense but are limited by the high overhead in practical usage. Moreover, in cloud environments, containers such as LXC now can be used as process and application-level virtualization to run multiple isolated systems on a single host. In this paper, we present a novel system called Memvisor to provide high availability memory mirroring. It is a software approach achieving flexible multi-granularity memory mirroring based on virtualization and binary translation. We can flexibly set memory areas to be mirrored or not from process level to the whole user mode applications. Then, all memory write instructions are duplicated. Data written to memory are synchronized to backup space in the instruction level. If memory failures happen, Memvisor will recover the data from the backup space. Compared with traditional software approaches, the instruction level synchronization lowers the probability of data loss and reduces the backup overhead. The results show that Memvisor outperforms the state-of-the-art software approaches even in the worst case.