Improving I/O Performance in Container with OverlayFS with Optimized Synchronization

Naoki Mizusawa
Electrical Eng. and Electronics
Kogakuin Univ. Graduate School
Tokyo, Japan
cm18039@ns.kogakuin.ac.jp

Yuya Seki
National Institute of Advanced Industrial Science and Technology
Tsukuba, Japan
yuya-seki@aist.go.jp

Jian Tao
Texas A&M Engineering Experiment Station and High Performance Research Computing
Texas A&M University
Texas, USA
jtao@tamu.edu

Saneyasu Yamaguchi
Electrical Eng. and Electronics
Kogakuin Univ. Graduate School
Tokyo, Japan
sane@cc.kogakuin.ac.jp

CCS CONCEPTS

•**Operating systems**

KEYWORDS

container, OverlayFS, Docker, copy-up

1. INTRODUCTION

Virtualization has been utilized for a variety of purposes, such as server consolidation for saving placing spaces and power consumption. Docker is a popular container-based virtualization system. It provides a virtualized environment with a small overhead. Because of its small overhead, we can expect that a highly consolidated virtualized environment, in which many containers run concurrently in a physical computer, can be achieved with Docker. Docker utilizes OverlayFS, which is a union mount filesystem, in order to present a filesystem to each container with a small overhead. In our previous work [1], we proposed a method for improving file access performance of OverlayFS by reducing the frequency of synchronization.

In this poster, we explain our proposed method and evaluate the I/O performance in a highly consolidated virtualized environment.

2. PROPOSED METHOD

This section describes our proposed method to improve I/O performance in Docker container processes [1]. OverlayFS synchronizes data in the cache memory and external storage by calling `vfsfsync` for each write operation that involves copy-up. We argued that synchronization on every write operation with copy-up was too pessimistic. To solve this problem, we proposed to reduce the frequency of `vfsfsync` calls. For example, this method reduces the frequency from once per one write operation to once per 50 write operations.

3. EVALUATION

In this poster, we evaluate the I/O performance in a highly consolidated environment with Docker containers using the OverlayFS improved by our proposed method.

We evaluated the writing performances in highly consolidated Docker containers using OverlayFS with and without our proposed method in the cases Copy-up occurs. The frequency of synchronization of the proposed OverlayFS was once per 50 write operations. We measured the I/O performance with FFSB (Flexible Filesystem Benchmark). This benchmark repeated to write a randomly selected file. For this benchmark, we created 4096 files in the lower directory on OverlayFS. These files were shared by all the containers. Their sizes were between 1KiB and 1GiB. The number of containers in a physical computer ranged from one to 600 and an FFSB process was executed in every container. For comparison, we measured the performance using DM-thin in the same setup. DM-thin does not occur Copy-up.

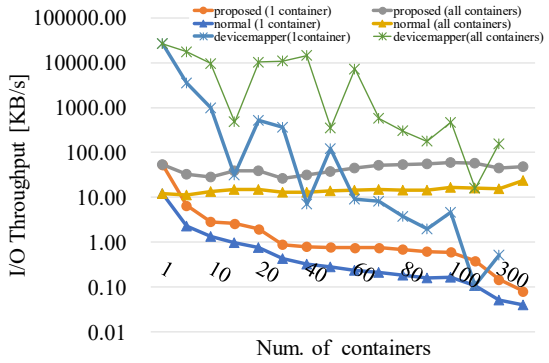

Fig. 1 I/O performance

Fig. 1 shows the results. From the figure, we can see that the I/O performance was significantly improved by our proposed method. In addition, the proposed OverlayFS could provide better than DM-thin even if the operations occurred copy-up in the cases of highly consolidated environments.

4. CONCLUSION

Our evaluation showed that the proposed OverlayFS outperforms the original OverlayFS and DM-thin in a highly consolidated container environment. We plan to evaluate our method in a variety of situations.

ACKNOWLEDGMENTS

This work was supported by JSPS KAKENHI Grant Numbers 17K00109, 18K11277. This work was supported by JST CREST Grant Number JPMJCR1503, Japan.

REFERENCES

[1] N. Mizusawa, J. Kon, Y. Seki, J. Tao, S. Yamaguchi, "Performance Improvement of File Operations on OverlayFS for Containers," 2018 IEEE International Conference on Smart Computing (SMARTCOMP), 2018, Vol. 1, pp. 297-302. doi: 10.1109/SMARTCOMP.2018.00019