ECI-Cache: A High-Endurance and Cost-Efficient I/O Caching Scheme for Virtualized Platforms

Saba Ahmadian
Sharif University of Technology
Tehran, Iran
ahmadian@ce.sharif.edu

Onur Mutlu
ETH Zürich
Switzerland
omutlu@ethz.ch

Hossein Asadi*
Sharif University of Technology
Tehran, Iran
asadi@sharif.edu

ABSTRACT

In recent years, high interest in using *Virtual Machines* (VMs) in data centers and cloud computing has significantly increased the demand for high-performance data storage systems. A straightforward approach to providing a high-performance storage system is using *Solid-State Drives* (SSDs). Inclusion of SSDs in storage systems, however, imposes significantly higher cost compared to *Hard Disk Drives* (HDDs). Recent studies suggest using SSDs as a caching layer for HDD-based storage subsystems in virtualized platforms. Such studies neglect to address the endurance and cost of SSDs, which can significantly affect the efficiency of I/O caching. Moreover, previous studies *only* configure the cache size to provide the required performance level for each VM, while neglecting other important parameters such as cache *write policy* and *request type*, which can adversely affect both performance-per-cost and endurance.

In this paper, we propose a new *high-Endurance and Cost-efficient I/O caching* (ECI-Cache) scheme for virtualized platforms in large-scale data centers, which improves both performance-per-cost and endurance of the SSD cache [1]. ECI-Cache dynamically assigns 1) an efficient cache size for each VM, to maximize the overall performance of the running VMs and 2) an effective write policy for each VM, to enhance the endurance and performance-per-cost of the storage subsystem.

To do so, we propose a new metric, called *Useful Reuse Distance* (URD), so as to minimize the allocated cache size for each VM while maintaining the performance of the VM. URD reduces the cache space by considering the request type in calculating the reuse distance of the workloads and eliminates the unnecessary write accesses (i.e., writes to a block without any further read access). ECI-Cache employs URD and, thus, allocates much smaller cache space for each VM compared to state-of-the-art cache partitioning schemes. ECI-Cache also efficiently decides what type of write policy is the best to use for each VM in the SSD cache. As a result, ECI-Cache is able to 1) maximize the performance-per-cost and 2) improve the endurance of the SSD cache.

We implemented ECI-Cache on an open-source hypervisor, QEMU (version 2.8.0), on the CentOS 7 operating system (kernel version 3.10.0-327). We evaluate our scheme on a real system where the hardware includes a 2U rack mount HP ProLiant DL380 Generation 5 (G5) server, 4x 146GB SAS 10K HP HDDs (in RAID-5 configuration), a 128GB Samsung 850 Pro SSD, 16GB DDR2 DRAM, and 8x 1.6GHz Intel(R) Xeon CPUs. We run more than fifteen workloads from the SNIA MSR traces on the proposed platform on 16 concurrent VMs. Experimental results we obtain reveal that ECI-Cache 1) improves performance by 17% and performance-per-cost by 30%, 2) reduces the number of committed writes to the SSD by 65%, all compared to the state-of-the-art dynamic cache partitioning scheme. Due to the large reduction in number of committed writes to the SSD, ECI-Cache also greatly enhances the endurance of the storage subsystem. We conclude that ECI-Cache is an effective method for managing the SSD cache in virtualized platforms.

For more information on ECI-Cache as well as the detailed experimental studies we conduct on the effect of 1) the URD metric and 2) the write policy of the SSD cache on system performance and endurance, please refer to our paper [1]. We hope that our results and the idea of the ECI-Cache inspire more research into new reuse distance metrics and new I/O caching mechanisms that take into account endurance and per-VM SSD access characteristics.

*Corresponding Author

SIGMETRICS'18 Abstracts, June 18–22, 2018, Irvine, CA, USA

ACM ISBN 978-1-4503-5846-0/18/06.
https://doi.org/10.1145/3219617.3219646

ACM Reference Format:
Saba Ahmadian, Onur Mutlu, and Hossein Asadi. 2018. ECI-Cache: A High-Endurance and Cost-Efficient I/O Caching Scheme for Virtualized Platforms. In *SIGMETRICS'18 Abstracts: ACM SIGMETRICS International Conference on Measurement & Modeling of Computer Systems Abstracts, June 18–22, 2018, Irvine, CA, USA.* ACM, New York, NY, USA, Article 4, 1 page. https://doi.org/10.1145/3219617.3219646

REFERENCES

[1] Saba Ahmadian, Onur Mutlu, and Hossein Asadi. 2018. ECI-Cache: A High-Endurance and Cost-Efficient I/O Caching Scheme for Virtualized Platforms. *Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS)* 2, 1 (2018), 9.