SMIO: I/O Similarity Aware Virtual Machine Management in Virtual Desktop Environments

Min Li†, Sushil Mantri†, Pin Zhou‡, Ali R. Butt†
† Dept. of Computer Science, Virginia Tech; ‡ IBM Almaden Research Center
{limin,sushil,butta}@cs.vt.edu, pinzhou@us.ibm.com

ABSTRACT

Virtual desktop environments (VDEs) typically employ shared storage with the goal to benefit from the I/O optimizations therein. However, the Virtual Machine (VM) management that enables VDEs is typically agnostic of VM I/O characteristics. Thus, the opportunities for I/O reduction techniques, e.g., deduplication, which improve storage scalability and efficiency are constrained. We present SMIO, a VM management system that detects VM I/O similarities, and places the VMs such that the effectiveness of I/O reduction techniques is enhanced—as much as 4.9× compared to the standard approach.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems — *distributed applications*
General Terms: Algorithms, Performance, Evaluation.
Keywords: virtual machine placement and migration

1. INTRODUCTION

Virtual desktop environments (VDEs) are becoming popular due to their lower cost, high performance and scalability. Shared storage is a key component of VDEs, and I/O optimizations are considered crucial to sustaining high scalability. In this context, researchers have observed large amounts of duplicate IOPS across VDEs. This is because the Virtual Machine (VM) images are usually created using the same golden image and deploy a similar set of applications such as anti-virus software. A number of I/O reduction techniques such as dedup-box, Atlantis ILIO, Capo [FAST'11], and Seacache [MASCOTS'12], have been proposed to identify and remove such duplicated I/O load from the shared storage system to improve efficiency. The effectiveness of these techniques depends on the amount of duplicated data accessed by the VMs running on a physical host. The goal of this work is to avoid sub-optimal VM placement and colocate similar VMs on the same hosts, which increases similar accesses on the hosts, consequently creating higher opportunity for I/O reduction.

2. DESIGN

We propose SMIO, a VM placement technique that considers I/O similarity among VMs to improve the efficiency of I/O reduction techniques by locating VMs with similar I/Os on the same physical host.

Figure 1 shows the architecture of SMIO. SMIO: (1) detects I/O similarity among different VMs; (2) utilizes hierarchical clustering to produce a new I/O-similarity-aware VM placement scheme periodically; and (3) migrates the VMs when benefits of similar VM consolidation outweigh migration cost. We employ sampling instead of continuous I/O monitoring to control the network and I/O overhead and ensure scalability to thousands of VMs. We also adopt a hybrid hierarchical clustering technique, where computation of cost-benefit analysis of VM migration is offloaded to the individual physical hosts on which the VMs are running, and the results from the hosts are combined at a centralized manager to determine global VM migration scheme.

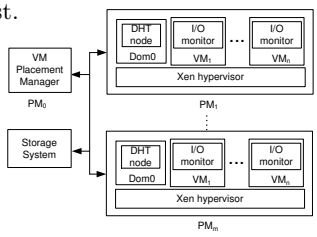

Figure 1: SMIO architecture.

3. EVALUATION

We evaluated the effectiveness of SMIO through trace driven simulations. Figure 2 shows that SMIO effectively detects similarities between different set of VMs and reduces the I/O bandwidth consumption by as much as 4.9× for read path and 2.1% for write path compared to state-of-the-art FFD [VEE'09] that manages VMs without considering I/O similarities.

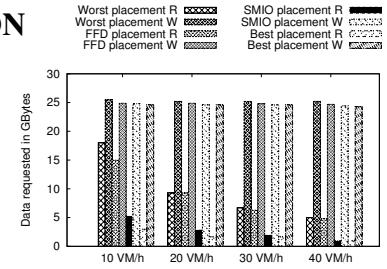

Figure 2: I/O bandwidth consumption under training center trace.

4. CONCLUSION

By incorporating I/O similarity information in VM management decision, SMIO is able to improve the effectiveness of I/O reduction techniques. Moreover, SMIO paves the way for designing innovative holistic I/O optimizations, which further improves scalability and storage efficiency.

Acknowledgment: This work is supported in part by the NSF under Grants CCF-0746832, CNS-1016793, and CNS-1016408.

ICS'13, June 10–14, 2013, Eugene, Oregon, USA.
ACM 978-1-4503-2130-3/13/06.