

Quantifying and Improving I/O Predictability in Virtualized Systems

Cheng Li, Ínigo Goiri, Abhishek Bhattacharjee
Ricardo Bianchini, Thu D. Nguyen
{chenglii, goiri, abhib, ricardob, tdnguyen}@cs.rutgers.edu
Department of Computer Science
Rutgers University, Piscataway, NJ 08854

Abstract

Virtualization enables the consolidation of virtual machines (VMs) to increase the utilization of physical servers in Infrastructure-as-a-Service (IaaS) cloud providers. Unfortunately, our quantification of storage I/O performance across a range of workloads, virtual machine monitor (VMM) architectures, approaches to storage virtualization, and storage devices shows widespread performance *unpredictability* in the face of consolidation. Surprisingly, the use of solid-state drives (SSDs) can exacerbate the problem. Since many users may desire consistent performance, we argue that IaaS cloud providers should provide a class of predictable-performance service in addition to their existing (predictability-oblivious) services. Thus, we propose and evaluate VirtualFence, a storage system that provides predictable performance for this new class of service. VirtualFence uses three main techniques: (1) non-work-conserving time-division I/O scheduling, (2) a small SSD cache in front of a much larger hard disk drive (HDD), and (3) space-partitioning of both the SSD cache and the HDD. Our evaluation of a prototype VirtualFence implemented in the Xen VMM shows that VirtualFence improves predictability significantly. More fundamentally, our evaluation illustrates the tradeoff between predictability and performance. We conclude that current VMMs are far from providing predictability. Systems like VirtualFence can remedy this problem, while allowing the cloud provider to select an appropriate compromise between performance and predictability.