Volume 2, Issue 6, June 2012 ISSN: 2277 128X
International Journal of Advanced Research in Computer Science and Software Engineering
Research Paper
Available online at: www.ijarcsse.com

Improving Throughput in Cloud Storage System

Arpit J. Kuche*, Mayur C. Patel, Kaushal P. Patel
Department Of Information Technology
Sant Gadgebaba Amravati University, India
arpit.msgs@gmail.com

***Abstract*—Because the cloud serves many workloads concurrently, its disk access pattern is highly random and heterogeneous. In addition, because various virtual machines access to files respectively, metadata utilization and small write requests are increased. In order to build a system for these patters, we should analyse the modern techniques used at cloud system. First, we show that a SATA controller has enough processing capability to serve six disks without performance degradation. Motivated by this experiment, we compare three disk configurations by executing benchmark applications. From the result, we observe that accessing independent disks directly provides better performance than RAID-0 and RAID-5.**

***Keywords*— component; Storage architecture; RAID; SATA; Cloud computing; Performance measurement; NFS; iSCSI;**

I. INTRODUCTION

Cloud computing services have been rapidly growing in popularity in many areas due to its flexible pricing model and ease of use. As the cloud plays an important role in various services the cloud providers are faced with new challenges such as reduction of power consumption and guaranteeing service level agreements (SLAs). In order to improve performance and quality of service, there has been a lot of research about consolidation of compute nodes based on virtualization technologies. However, the bottleneck of the present cloud computing services is the storage system. Therefore, it is essential to improve the performance of the cloud storage system for better efficiency and throughput.

In these cloud storage systems, many workloads are conducted concurrently. Because of this, access patterns to disk are congested and cause an adverse effect to execution time. Also if each workload approaches a large number of files, the number of access to metadata is increased. Finally, small writes are accepted to disk frequently. It is directly related with performance.

II. SATA CONTROLLER TEST

We performed an experiment of a SATA controller to measure its performance with a varied number of disks. Our test only reads data from disks sequentially in order to measure the maximum throughput of the SATA controller. We performed the test with increasing the number of readers from one to six. We set each disk to be accessed by only one reader at a time.

A. Experimental Setup

Our experiments use a server with the following specifications.

- Four AMD 12-core Opteron Processor 6168 CPU (total 48 cores)
- 128GB RAM
- Six Western Digital 1TB HDD, 7200RPM
- ATI SB700/SB800 SATA Controller
- CentOS 6.0 (Linux 2.6.32).

B. Experimental Results

Fig. 1 shows the result for the test. As shown in the graph, the performance degradation of each disk is not large even with an increase of active disks and workloads except for Disk 6 in which OS is installed. It means that there is little contention in a SATA controller, and it has enough processing capability for serving six SATA disks.

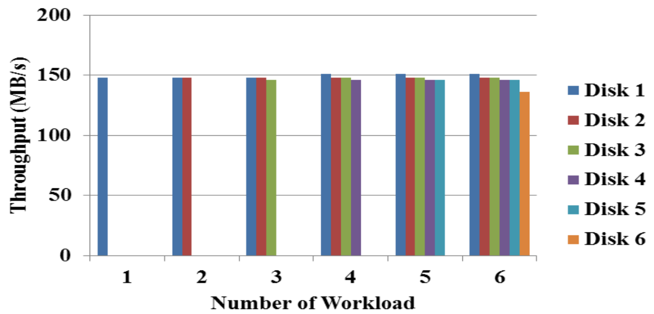

Fig. 1 Disk I/O performance test result in a server with a SATA controller of ATI SB700/SB800 for AMD Opteron. Although the number of active disks is increased from one to six, the performance degradation is not large.

III. NO-RAID AND RAID

In this work, we present performance evaluation on various storage architectures; it is motivated by the above experiment. Because the cloud serves many workloads concurrently, its disk access pattern is highly random and heterogeneous. RAID was designed for increased reliability and performance when there is small number of concurrently active users, such as supercomputer system. Therefore, accessing independent disks directly can yield better performance than RAID in some situations.

As mentioned before, a SATA controller can serve six SATA disks without performance degradation due to interference among them. In this Section, we compare the performance of several disk configurations by using the PostMark benchmark application.

A. Experimental Setup

We performed an experiment in same machine when used in Section 1.

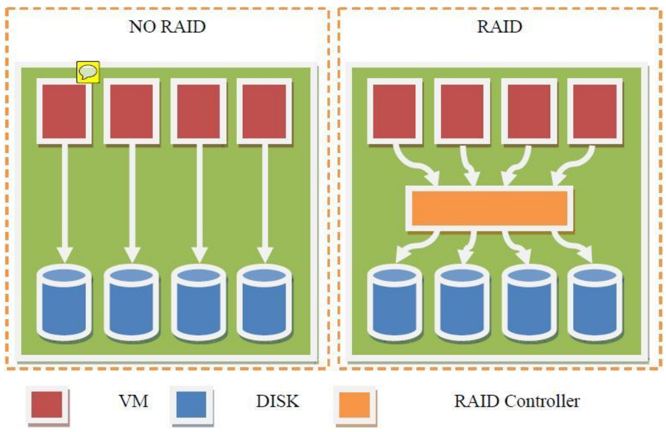

Fig. 2 Architectural comparison between No-RAID and RAID. In No-RAID, each workload can access disk directly. However, in RAID environments, all I/O requests and responses should pass the RAID controller.

In all experiments, Postmark is configured as below.

- Number of files: 100,000
- Number of transactions: 200,000
- File size: 100,000 bytes

We compared 3 different storage configurations using five disks as follows.

- No-RAID (Each disk serves different workload.)
- RAID-0
- RAID-5

The above two RAIDs are constructed by software. Fig.2 illustrates an architectural comparison between No-RAID and RAID. In No-RAID, each workload can access disk directly, and each disk can process request independently without interference by other disks. However, in RAID environments, all I/O requests and responses should pass the RAID controller. In I/O-intensive situation, the RAID controller can be a performance bottleneck. Especially, in RAID-5, each disk cannot serve request independently. In the following section, we present experimental results about these three disk configurations.

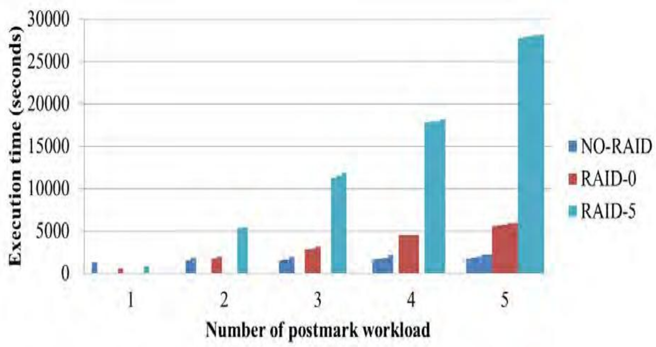

Fig. 3 Performance test results for three disk configurations. No-RAID shows the best performance except the first case.

B. Experimental Results

As you can see Fig. 3, there are considerable differences between No-RAID and two RAIDs. When just one postmark is run, the two RAID environments show better performance than No-RAID. However, as the number of active PostMark exceeds one, RAID-5 shows poor performance. This is because PostMark transactions are made up of a lot of random small write requests. And it is well-known that RAID-5 is vulnerable to numerous small writes.

Not as much as poor like RAID-5, RAID-0 also shows increasingly poor performance depending on number of Postmark workloads. On the other hand, there was almost no slowdown in No-RAID system. So, it can be a reasonable choice to use independent multiple physical disks when running several I/O-intensive workloads simultaneously.

In real world, although RAID-0 has the best performance among RAID series, it is seldom used because crash risk is too big. Because of that reason, it is normal that constructing storage system to RAID-5 or RAID 1+0. It means that using disks independently can be much useful in real system frequently running disk-intensive jobs.

IV. STORAGE PROTOCOL TEST

NFS and iSCSI are two network storage protocols used mainly when cloud provider constructs cloud systems. First, NFS has used often because it is easy to configure and use. When specific folder was opened by NFS server, clients just mount it to any location of local. So, clients don't know real file system in NFS server. In this system, if write operation occurred, NFS client cached them and transferred to server when it reached to specific condition. On the other hand, iSCSI that made by existing SCSI protocol to useable in network is hard to configure but has better performance and security. When using iSCSI, target allocates any size of disk space and services them by iSCSI target daemon. And initiator can recognize volumes as a new block device. Unlike NFS, initiator can install the file system to recognized iSCIS target volumes freely and mount to them on specific point. iSCSI also caches requests when it is write operation. This cache policy depends on installed file system.

Likewise pervious Sections, these protocols also are loaded heavy small random I/O workloads because many virtual machines access to them. So we can find which protocol is suitable this patterns by processing experiment using small random I/O benchmark.(Postmark)

A. Experiment Setup

Our host machine's specification is as below.

- Intel Core 2 Quad CPU Q9550 2.83GHz
- 8GB RAM
- Xen 4.0.1, Linux Kernel 2.6.32.25

Our storage server's specification is as below.

- Four 2TB SATA 5900RPM Hard Disk
- RAID-0

B. Benchmark

Likewise Section 2, we used PostMark as benchmark. Because it has characteristics like small random I/O pattern. PostMark divided 3 parts. First, it creates many folders and files when used processing time. Second, it append to created files and repeat it as much as the number of transactions defined by options. In this time, most operations are small random write requests. After that, finally it deletes all the files that used in processing time.

C. Experimental Results

First, we acquired result that PostMark runtime on NFS is 1.5 times slower than PostMark runtime on iSCSI. The reason for the above result is difference of the number of write operations. As you can Fig. 4, the number of write operations on NFS is almost 1.7 times more than the number of write operations on iSCSI. We can think that this difference comes from cache effect in I/O path. If many append operations merged in the caches, fewer operations are passed to the disk. We can find that cache effect is not negligible through another experiment. (Fig.5) On the Same situation, we just performed experiment while increasing the write buffer size of NFS.

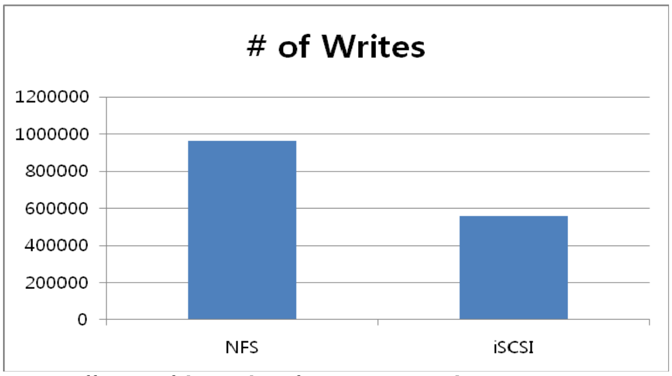

Fig. 4 Difference of the number of writes in NFS and iSCSI

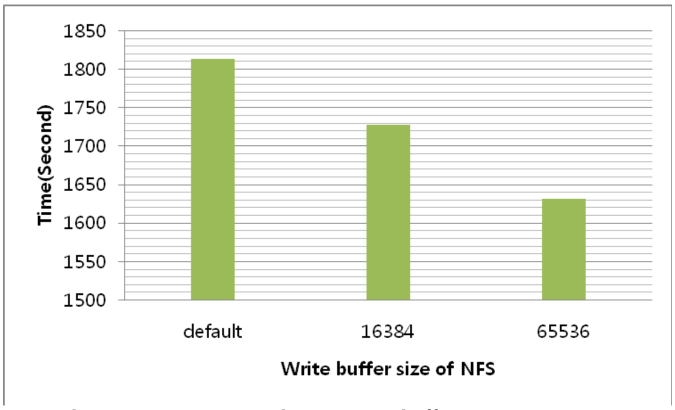

Fig. 5 Change in execution time by NFS write buffer size

V. Conclusion and Future Work

In this paper, we progressed three experiments to find whether modern techniques are adaptable to cloud storage systems. Firstly, we performed the experiment that SATA controller test. Through this experiment we observed that a SATA Controller can serve multiple disks without performance degradation due to interference among disks. And in the second experiment, we got the result that RAID system's performance is worse than No-RAID system when executing multiple disk-intensive workloads. Through these results, RAID architecture is not adequate to cloud computing environments. Although RAID showed bad performance in multiple workloads test, it still has good characteristics like abstraction layer. Lastly, we performed the experiment that which network protocol is adaptable to cloud storage systems. Using PostMark benchmark, we got the result that iSCSI showed better performance than NFS and cache effect is very effective to merge the write requests. It can say that iSCSI is more adaptable than NFS to cloud storage systems. Although it is used frequently in cloud systems, it is not just for cloud systems but for any other systems. So if new protocol only for cloud storage is created, it can show better performance than any other protocols.

So our next plan is developing the storage system that can offer independent use of disk and abstraction layer like RAID. And we will modify the existing network storage protocol code to much suitable for cloud storage systems.

References

[1] R. Buyya, C.S. Yeo, Venugopal, *Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities*, 10th IEEE Int. Conference on High Performance Computing and Communications, HPPC 2008, Dalian, China, Sept. 2008.

[2] Vouk, M. A. *Cloud computing Issues, research and implementations*, 30th International Conference on Information Technology Interfaces(ITI 2008(, Cavtat/Dubrovnik, Croatia, June 2008, 31-40.

[3] Johan Bresnahan, David LaBissoniere, Time Freeman, and Kate Keahey, *Cumulus: An Open Source Storage Cloud for Science*, San Jose, CA, 2011.

[4] Jeffrey Shafer, *I/O Virtualization Bottlenecks in Cloud Computing Today*, Proceedings of the Second Workshop on I/O Virtualization Pittsburgh, PA, USA, Mar. 2010.
[5] L. Tang, J. Mars, N. Vachharajani, and R. H. M. L. Soffa, *The impact of memory subsystem resource sharing on datacenter applications,* In Proceedings of ISCA '11, San Jose, California, USA, 2011.
[6] Y. Koh, R. Knauerhase, P. Brett, M. B. andZhihua Wen, and C. Pu, *An analysis of performance interference effects in virtual environments,* In In Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), 2007.
[7] Katcher, J, PostMark: A New File System Benchmark, System (3022), pp 1—8, 1997.
[8] W. Wang, H. Yeo, Y. Zhu, and T. Chong. "Design and development of Ethernet-based storage area network protocol". In Proceedings of the 12th IEEE International Conference on Networks (ICON), 2004.
[9] S. Sivathanu, L. Liu, M. Yiduo, and X. Pu. "Storage management in virtualized cloud environment." Proc. IEEE 3rd Int. Conf. On Cloud Computing, pp. 204-211, 2010.
[10] J. Satran, K. Meth, C. Sapuntzakis, M. Chadalapaka, and E. Zeidner. Internet small computer systems interface(iSCSI). RFC 3720 (Proposed Standard), April 2004.
[11] S. Shepler, B. Callaghan, D. Robinson, R. Thurlow, C. Beame, M. Eisler, and D. Noveck. NFS version 4 protocol. RFC 3010, Network Working Group, December 2000.

