How to Model Hard Drive Performance in the Cloud

Technology Paper

Written by
Dan Lingenfelter, PhD

People are creating and consuming data at such a rapidly growing rate that it challenges providers of cloud storage to manage increasingly large amounts of storage. A key part of efficient, low-cost cloud storage is the right hard disk drives (HDD). Because different HDDs perform better under different workloads, cloud storage providers can benefit from modeling HDD performance to determine which performance characteristics are best for various elements of a cloud storage provider's environment. Modeling can also be used to compare different HDDs to see which ones offer the best performance under a given workload. To obtain accurate estimates, cloud storage providers must have an understanding of the workloads their drives will handle and a set of basic performance parameters for each drive to be modeled.

Percentage of Random and Sequential Drive Transfers

To accurately model performance, determine the percentage of drive transfers that are random and sequential in a workload as well as the balance between read and write operations.

Sequential transfers occur when the data being accessed is located at adjacent addresses. Transfers that are not sequential and start at random addresses are called random. Workloads generated by an email database frequently have a large random component because messages are small, span few addresses and are often retrieved randomly. However, a backup application likely generates a mostly sequential workload because data is read and written in large increments that span many addresses.

How to Model Hard Drive Performance in the Cloud

Random Performance

When modeling random performance on an HDD, estimates of the following are required:

- Time for the head to seek from one location to another
- Time it takes the data to rotate to the head's location
- Efficiency of the native command queuing feature (if applicable) that reorders requests so they can be serviced more efficiently

An HDD's throughput in a random workload is often limited by the drive's seek time, rotations per minute (RPM) and the native command queuing (NCQ) configuration.

Sequential Performance

When modeling sequential performance on an HDD, estimates of the following are required:

- The number of usable bits that pass beneath the drive's head per second
- The drive's RPM
- The maximum linear density in bits per inch
- Platter radius

These parameters are usually available on an HDD's datasheet.

HDDs are designed to stay at one rotational speed. Because of this, the distance covered by bits stored in the outer platter in a second is higher than on inner tracks. For example, on a 3.5-inch platter in an HDD like the Seagate® Desktop HDD, the track on the outer edge is 10 inches long, while a track on the innermost edge is 2.5 inches. If a platter spins at 7200 RPM, the outer edge travels at a velocity of 67 miles per hour (MPH) whereas the inner edge travels at 17 MPH.

What this means is data stored at different places on a platter in a drive can be transferred faster or slower depending on location. Because transfer rate can vary by data location, it is best to average sequential performance over the entire partition being used.

Sequential vs. Random Workloads: Is Your Workload Sequential?

Determining what percentage of a workload is sequential is important since it affects HDD performance. A purely sequential workload yields much higher performance than a workload that is even slightly random.

During testing conducted by the Seagate Cloud Modeling and Data Analytics group, it was found in one case that the sequential transfer of a file on an ext4 file system was only 97.5% sequential at the drive level due to metadata transfers and requests from other system processes. The 97.5% sequential file transfer yielded significantly less performance than is expected for a purely sequential workload. The performance difference occurred because the time required to seek and rotate on the 2.5% of random transfers was very large compared to the time required to transfer data to or from the media. Figure 1 shows expected performance for a purely sequential workload and the performance measured in practice where the workload is less than 100% sequential due to metadata accesses.

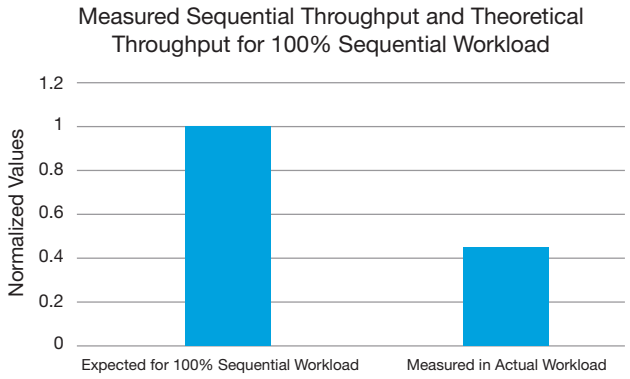


Figure 1. Measured Sequential Throughput and Theoretical Throughput for 100% Sequential Workload

Conclusion

Understanding HDD workload makes it easier to accurately model and estimate HDD performance. Providers of cloud storage can save time and money and increase performance by determining the kind of workloads they are likely to see in their environments and modeling HDDs accordingly. For more information about parameters for HDD modeling, refer to: Efficient Disk Drive Performance Model for Realistic Workloads in IEEE Magnetics.

www.seagate.com

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16–18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00

© 2014 Seagate Technology LLC. All rights reserved. Printed in USA. Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. Seagate reserves the right to change, without notice, product offerings or specifications.
TP661.1-1409US, September 2014