



Title: High velocity data, mass storage and access density.

Author: Bill Mottram

Date: November 10th, 2008

Abstract: Storage performance is traditionally defined by measures of data latency and bandwidth. However, in the case of scale-out, large volume performance it is necessary to look beyond these traditional measures and consider I/O performance and how easily compute resources can be scaled to keep performance in balance with capacity growth. Higher IOP drive higher access densities which improves the ability of storage solutions to manage peak query volumes and variable query complexity. A particularly critical metric in today's internet driven world.

While doing some research for a recent whitepaper I came across the term “high velocity data”. A term that resonated with me but caused me to wonder what the implications of high velocity data were for storage architectures.

The term referred to the speed at which critical business data flows in and out of an enterprises, how quickly data movement is assimilated into the data pool and how quickly data updates are available to support ongoing business operations. Consider Snapfish the on-line photo folks. Unless their data delivery infrastructure is very nimble and responsive their on-line customers will get disillusioned and switch to a competitor. This high level of responsiveness needs to be maintained despite a highly variable query volume which on a peak day can be 10 million print requests. A recent Information Week report talked about LGR Communications and their data warehousing project that supports a growing 5000+ user base that generates about 10 million queries daily, creating an update volume that touch billions of records. Clear illustrations of high

velocity data at work and where the survival of an enterprise is dependent on a storage infrastructure that cannot just meet the challenge effectively, but do so affordably.

But what measures the goodness of a storage solution when considering how it will handle high velocity data?

Obviously bandwidth is important and depending on the application could be the critical performance parameter. However, just as critical is data latency. This is the measure of how quickly data is made available for use following a request. The obvious goodness factor in this e-commerce, internet age is the speed at which user screens are refreshed following a query. Think of your own frustrations when you are trying to navigate a lethargic site, you do not stay there long. The determining variable for data latency is access density which establishes how this *wait time* behaves under the pressure of a high query hit rate. So what is access density?

Access Density is defined as the ratio of drive I/O over capacity (IOP/GB). If capacity doubles and performance doubles then access density remains unchanged. In April 2000, Fred Moore wrote in CTR *“that while the capacity of a disk drive has increased 6000 times since 1964 the raw performance, seek, latency and transfer rate has only increased by a factor of 8”*. A massive imbalance and points to the issue that scaling disks is more than just capacity and despite techniques to counter the problem such as larger cache and actuator level buffers the imbalance remains. In short as areal density has increased, more capacity sits under actuators and the access performance of these drives drop. This access performance hit extends to the subsystems that use these drives.

Unless there is someone with greater wisdom that can correct me, the answer when driving data latency improvements is to have less data under each actuator. The question is how?

One approach is to short stroke the drives. This means that the placement of highly active data is restricted to the outer bands of the disk. Good way to waste a lot of expensive disk unless policies are in place (and enabled) to harvest the inner bands for the not so high velocity data. Pillar Data is an example where onboard QoS policies determine where the data resides on the disk.

As more spindles increase I/O performance, another option would be to increase the number of spindles preferably in the same or better physical footprint. A large number of densely packaged spindles is something that the new generation of clustered, highly dense storage solutions enjoy. Today most solutions use 3.5” FF drives but by substituting them for a 2.5” FF, packaging density and spindle count can be increased

Veridictus Associates Inc

14 Via Sierra Grande, Manitou Springs, CO 80829 – (303) 588-6674

web: www.veridictuassociates.com

blog: www.storagetopics.com

without increasing footprint. Increasing system spindle count increases the access density ratio. Two companies are already successfully exploiting 2.5” drives in enterprise class storage solutions, Atrato and Xiotech and judging from the recent Seagate announcement for their Sawio 15k 2.5” drive HP and Dell are not far behind. For those green conscious folks note that these smaller form factor drives consume much less power, 70% less than their 3.5” bigger brother.

Back to my original question. Access density is a blend of I/O performance and capacity with possibilities inevitably tempered by cost. Each must be balanced to reach the compromise that will deliver the performance needed at a cost that is affordable. Scale-out storage offers the most likely opportunity for succeeding in this quest. Perhaps something to remember when considering a future storage developments or acquisitions.

High velocity data is the transactional lifeblood of commerce. So the ability of a storage subsystem to affordably manage large volumes of high velocity data means that the storage solution must be optimized to deliver a high access density that keeps data latency low which in turn will drive the ultimate measure of success, user satisfaction.

Veridictus Associates Inc

14 Via Sierra Grande, Manitou Springs, CO 80829 – (303) 588-6674

web: www.veridictuassociates.com

blog: www.storagetopics.com

Bio: Bill Mottram compliments his role the managing partner of Veridictus Associates Inc, a high technology marketing consultancy as analyst practitioner with Data Mobility. He has over 25 years product development, marketing and sales experience in the data storage and information technology industries. Experienced with Fortune 500 companies such as StorageTek, Compaq and HP and smaller, start-up enterprises including Pillar Data Systems and COPAN Systems he was the marketing executive and leader responsible for the development , introduction and “go-to-market activities for a number of innovative and highly successful data storage solutions. Although primarily a marketing professional Mottram has considerable expertise in technology management, corporate acquisition and public company funding.

A graduate of Paisley College of Technology (electronic engineering) and holds a Masters of Business Administration from Pepperdine University.

Mottram can be contacted at bill@veridictusassociates.com

Copyright © 2008 Veridictus Associates Inc. All Rights Reserved. Reproduction of this publication without prior written permission is forbidden. Veridictus Associates Inc believes the statements contained herein are based on accurate and reliable information. However, because information is provided to Veridictus Associates Inc from various sources, we cannot warrant that this publication is complete and error-free. Veridictus Associates Inc disclaims all implied warranties, including warranties of merchantability or fitness for a particular purpose. Veridictus Associates Inc shall have no liability for any direct, incidental, special, or consequential damages or lost profits. The opinions expressed herein are subject to change without notice.

Veridictus Associates Inc

14 Via Sierra Grande, Manitou Springs, CO 80829 – (303) 588-6674

web: www.veridictuassociates.com

blog: www.storagetopics.com