Scaling Capacity and Performance

Disk and Tape *Scale* Differently

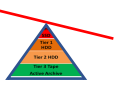

By:
Horison Information Strategies
Fred Moore, President
www.Horison.com

Addressing the amount of digital data created is *the* target for the storage industry. The world now produces nearly ~2.5 exabytes ($1x10^{18}$) of data a day and much of it will be preserved indefinitely with hopes of someday capitalizing on its potential value. Even with the amount of data created that is sensed, overwritten, discarded, but never stored longer than a few milliseconds, an incessant demand for storage capacity still exists across industries, governments, enterprises, and consumers. **Scalability** is the property of a system to handle a growing amount of work by adding resources to the system. For storage systems, the ideal strategy is to scale up (increase) capacity *and* performance together, however, for pricing considerations, more of the focus has historically been on scaling capacity first, then performance. This has led to the rapid growth of SSDs and NVM systems to address high performance storage requirements.

Disk and Tape Capacity Scaling Considerations - Adding capacity for disk and tape however is quite different. Adding disk drive capacity is straightforward but it gets costly. For disk, whenever additional disk capacity is needed a new drive and HDA (Head Disk Assembly – comprising R/W heads, motors, fans, disk platters, enclosure and actuators) has to be installed. This requires additional cost, energy and footprint. **The ratio of a disk capacity increase to disk drive increase is 1:1.**

For data centers with robotic tape libraries, adding tape capacity normally is accomplished by simply adding additional cartridges to the library. For tape libraries, the ratio of total cartridges to attached drives (total carts/total drives) is *the* key metric for determining tape storage cost and throughput. This ratio ranges between 100 to 300 (or more for deep archives) cartridges per attached tape drive with an average of ~200 cartridges per drive. Scaling tape capacity is easier than scaling HDD capacity and yields huge cost savings as archives and long-term storage demand soars. In the example below, 20 tape drives in a robotic library support the same capacity as 4,000 disk drives. **The ratio of tape cartridge capacity increases to a tape drive increase is ~200:1.**

Disk and Tape Scaling

- **Disk scales by adding another drive when capacity is needed.**
- **Medium to large scale robotic tape libraries average 100 - 300+ cartridges per each attached drive.**
- **Assume 200 cartridges per drive as the library average.**
- **Assume cartridge capacity is 20 TB and disk capacity is 20 TB.**
- **Tape libraries typically scales capacity by adding one tape drive for every 200 cartridges added.**

Capacity (PB)	Tape Cartridges (@20TB)	Library Attached Tape Drives	Disk Drives (@20TB)
10	500	2.5(3)	500
20	1,000	5	1,000
30	1,500	7.5(8)	1,500
40	2,000	10	2,000
80	4,000	20	4,000

Source: Horison, Inc.

Disk and Tape Performance Scaling Considerations – The total throughput (performance – data rate) of a tape library increases, along with the cost, as more drives are attached. A tape library with a ratio of 100 cartridges per tape drive would offer higher throughput than one averaging 300 cartridges per drive. For tape drive performance, the data rate of tape drives steadily increases as capacity increases. Libraries with cold data will use fewer cartridges per drive than active archive libraries will. For the LTO family of drives, LTO-1 had a data rate of 15 MB/sec and the latest version, LTO-8 has a data rate of 360 MB/sec. - 24 times greater. The enterprise tape drive family has increased data rates from 100 MB/sec with the TS1120 to 400 MB/sec for the latest TS1160 drive.

Future performance challenges for disk drives (HDDs) are mounting. Access density is the measure of IOPS per Terabyte (see chart below). HDDs are increasing in capacity but *not* performance as the IOPS (I/Os per Second) improvements for HDDs have basically leveled off resulting in a steady access density decline.

HDD Scalability Challenges

Performance Gains are Negligible for HDDs

$\int(\frac{x}{y})$ Access Density = $\frac{\text{IOPS}}{\text{TB}}$

IOPS @ 10ms	HDD Cap. TB	Access Density
100	1.0	100
100	4.0	25
100	8.0	12.5
100	16.0	6.25

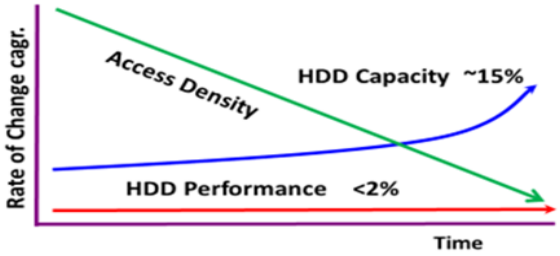

- HDD Performance (Speed) ***Not*** Scaling With HDD Capacity Growth or Server Speed
- Future HDD Performance Gains are Minimal - <u>if Any</u>
- Access Density Will Continue to Decline as HDD Capacity Increases
- Creates Additional Demand For SSD/NVM (Tier 0) Systems
- Results in HDD Capacity Reductions to Maintain Performance (Short Stroke) – Or Less Active Files

Source: Horison, Inc.

The potential for more concurrently active files increases as HDD capacity grows, and contention for the single actuator arm increases causing erratic response times. HDD areal density, which is directly related to device capacity, is currently progressing at ~16% annually, about half the rate of tape technology.

Conclusion

For many data centers, non-stop data growth can mean adding storage capacity on a daily or weekly basis. Neither disk nor tape scales capacity and performance the same. Adding disk capacity means adding another drive while a tape library can typically add 100-300 cartridges or more before another drive is required. For example, using 20TB device capacities, for disk you can add one drive for every 20TB - for tape you typically add one drive for every 4PB. Each new disk drive generation increases capacity but the IOPS have essentially levelled off causing the access density to steadily decline. For tape, each new drive generation increases capacity <u>*and*</u> throughput. Data centers need to scale disk and tape capacity to meet future demands. Moving lower activity data from disk to tape will free costly disk space, therefore reducing the number of disk drives needed for future growth while improving disk drive performance and shrinking the backup window. The overall economic benefits of correctly scaling tape and disk capacity are compelling particularly as the amount of data gets larger.