

Storage Devices

KALADHAR VORUGANTI

Network Appliance, Sunnyvale, CA, USA

5

Definition

One of the goals of database, file and block storage systems is to store data persistently. There are many different types of persistent storage devices technologies such as disks, tapes, DVDs, and Flash. The focus of this write-up is on the design trade-offs, from a usability standpoint, between these different types of persistent storage devices and not on the component details of these different technologies.

10
15

Historical Background

From a historical standpoint, tapes were the first type of persistent storage followed by disks, CDs, DVDs, and Flash. Newer types of memory technologies such as PRAM and MRAM are still in their infant stages. These newer non-volatile memory technologies promise DRAM access speeds and packaging densities, but these technologies are still too expensive with respect to cost/gigabyte.

20
25

Foundations

- Tapes/Tape Libraries: Tape readers/tape head, tape library, tape robot, and tape cartridge are the key components of a tape subsystem. Tapes provide the best storage packaging density in comparison to other types of persistent storage devices. Tapes do not provide random access to storage. Data on tapes can be stored either in compressed or uncompressed format. Unlike disks, tape cartridges can be easily transported between sites. Most organizations typically migrate data from older tape cartridges to newer tape cartridges every 5 years to prevent data loss due to material degradation. One can employ disk based caches in front of tape subsystems in order to allow for tapes to handle bursty traffic. Tapes that provide Write-Once, Read Many (WORM) characteristics are also available. WORM tapes are useful in data compliance environments where regulations warrant guarantees that a piece of data has not been altered. DLT and LTO are currently the two dominant tape technologies in the market. Technology-wise both these standards have minor differences. Finally, from a pure media cost standpoint, tapes are less expensive (cost per gigabyte) than disks and other forms of persistent media.

30
35
40
45
50

- Disks/Storage Controllers/NAS Boxes: Disks are the most widely used form of persistent storage media. Disks are typically accessed by enterprise level

55

applications when they are packaged as part of the processing server box (direct attached storage model), or are part of a network attached storage box (NAS) and accessed via NAS protocols or, are packaged as part of a storage controller box and accessed via storage area network protocols (SAN). The current trend is for protocol consolidation, where the same storage controller provides support for both SAN and NAS protocols. Typically, the size of the storage controllers can vary from a few terabytes to hundreds of terabytes (refrigerator sized storage controllers). A storage controller typically consists of redundant processors, protocol processing network cards, and RAID processing adapter cards. The disks are connected to each other via either arbitrated loop or switched networks. Storage controllers also contain multi-gigabyte volatile caches. Disks are also packaged as part of laptops.

60

65

70

75

80

85

- DVD/Juke Boxes: DVDs and CDs are optical storage media that provide random access and WORM capabilities. Only recently, the multiple erase capacity of an individual CD, or DVD was less than the capacity of a single disk drive or tape cartridge DVDs can store more data than a CD and a high definition DVD can store more data than a DVD. There are numerous competing standards for CDs, DVDs and high definition DVDs, however, format agnostic DVD players and DVD writers are emerging. Usage of DVDs is more prominent in the consumer space rather than in the enterprise space. A juke box system allows one to access a library of CDs or DVDs. DVDs have slower access speeds than most types of disks.

90

95

100

- Flash/SSDs/Hybrid Disks: Flash is memory technology that has non-volatile characteristics. Flash memory has slower read times than DRAM. Moreover, it has much slower write times than DRAM.

105

110

One has to perform an erase operation before one can re-use a flash memory location. One can only perform a limited number of erase operations. Thus, the number of write operations determines the Flash memory life. SLC and MLC are the two different NAND flash technologies. SLC can be erased a greater number of times, and it has faster access times than MLC based

flash. NAND flash has faster write and erase times than NOR flash. NOR flash has faster read times than NAND flash. NAND flash is used to store large amounts of data, whereas NOR flash is used to store executable code.

115
120
125
130
People are using MLC flash in cameras and digital gadgets, and are using SLC flash as part of solid state disks (SSDs). SSDs provide block level access interface (SCSI), and they contain a controller that performs flash wear leveling and block allocation. Hybrid disks that contain a combination of disks and Flash are emerging. Hybrid disks provide a Flash cache in front of the disk media. One typically can store meta-data or recently used data in the flash portion of hybrid disks to save on power consumption. That is, one does not have to spin up the disk. Flash storage provides much better random access speeds than disk based storage.

Key Applications

135
140
145
Tapes are being used primarily for archival purposes because they provide good sequential read/write times. Disks are the media of choice for most on-line applications. Optical media (CDs, DVDs) are popular in the consumer electronic space. Flash based SSDs are popular for those workloads that exhibit random IOs. Disks are being used in laptops, desktops and storage servers (SANs, NAS, DAS). Tape based WORM media and content addressable based disk storage are providing WORM media capabilities in tape and disk technologies, and thus, these technologies can be used to also store compliance/regulatory data.

(Abridged)

Recommended Reading

150
155
1. Anderson D., Dykes J., and Riedel E. More than an interface- SCSI versus ATA. In Proc. Second Annual Conf. on FAST. San Francisco, CA, 2003.
2. Toigo J. Holy Grail of Network Storage Management. Prentice Hall, Englewood Cliffs, NJ, 2003.
3. Voruganti K., Menon J., and Gopisetty S. Land below a DBMS. SIGMOD Rec., 33(1):64–70, 2004.

Answer the following questions:

- 1) Name some of the persistent storage devices mentioned in the text.
- 2) What are the advantages and disadvantages of tapes?
- 3) Where are WORM tapes useful?
- 4) Name the two dominant tape technologies that are currently on the market.
- 5) What is the most widely used form of persistent storage medium?
- 6) What are the protocols which are currently in use for accessing disks?
- 7) What spin-down modes are used to save on power consumption?
- 8) What are the characteristics of some of the optical storage media?
- 9) What is flash technology?
- 10) Where is flash memory used?

Mark the following statements as *true* or *false*:

- 1) Portable and inexpensive to purchase, tapes are often used for backing up or archiving data.
- 2) Flash is memory technology that has volatile characteristics.
- 3) WORM disks can be written only once.
- 4) NAND has significantly lower storage capacity than NOR.
- 5) MLC is used in solid state drives (SSDs) and SLC is used in consumer applications like cameras, media players, cell phones, etc.
- 6) An SSD (solid-state drive or solid-state disk) is a storage device that stores persistent data on solid-state flash memory.
- 7) In hybrid disks the flash memory handles the data most frequently written to or retrieved from storage.

Match the following terms with their definitions:

- 1) A library
 - 2) A disk cache
 - 3) Fibre channel
 - 4) A storage area network (SAN)
 - 5) DLT (digital linear tape)
 - 6) A jukebox
-
- a) A mechanism for improving the time it takes to read from or write to a hard disk
 - b) A unit in which optical disk drives are mounted
 - c) A high-speed special-purpose network (or sub-network) that interconnects different kinds of data storage devices
 - d) A collection of physical storage media such as tapes or disks and a way to access them
 - e) A form of magnetic tape and drive system used for computer data storage and archiving
 - f) A type of high speed interconnection

Vocabulary

access time ['æk.ses] [taim] - přístupová doba
bursty traffic [bə:sti] [træfik] - shlukový provoz, shluky dat
cache [kæʃ] - skryš, úkryt
competing [kəm'pi:t] - konkurenční
data compliance ['deɪ.tə] [tə] - ochrana dat před přepsáním
[kəm'plai.ənts] - ochrana dat před přepsáním
wear leveling [weər] [wer] ['lev.əl] - strategie zápisu a mazání (vyrovnání opotřebení)
focus on st ['fəʊ.kəs] [fou-] - zaměření
format agnostic ['fɔ:.mæt] [fɔ:r-] - jakýkoli formát
[æg'nɒs.tɪk] [-'nɑ:.stɪk] - přehrávací
gadget ['gædʒ.ɪt] - přístroj, zařízení
goal [gəʊl] [gou] - cíl
handle ['hæn.dl] - zvládnout, vypořádat se
high definition [haɪ] [def.r'nɪf.ən] - vysoké rozlišení
idle ['aɪ.dl] - nečinný
life [laɪf] - životnost
marked [mɑ:kt] [mɑ:kt] - výrazný
non-volatile [non'vɒl.ə.taɪl] [non'vɑ:l.ə.tə] - stálý, stabilní
numerous ['nju:z.mə.rəs] [nu:-] - početný, četný
persistent [pə'sɪs.tənt] [pə:-] - stálý, trvalý
read time [red] [taim] - přístupová doba pro čtení
solid state disk (SSD) ['sɒl.ɪd] [sa:ɪd] [steɪt] [dɪsk] - pevný disk na bázi paměťových čipů
standpoint ['stænd.pɔɪnt] - stanovisko, hledisko
to access st ['æk.ses] - dostat se k
to allocate ['æl.ə.keɪt] - přidělit
to alter ['ɒl.təɪ] [a:l.tə] - změnit
to determine [dɪ'tɜ:.mɪn] [tɜ:-] - určovat, udávat
to emerge [ɪ'mɜ:dʒ] [mɜ:dʒ] - objevit se
to package ['pæk.ɪdʒ] - zabalit
to spin down [spɪn] [daʊn] - zpomalit otáčení

to spin up [spɪn] [ʌp] - zrychlit otáčení
to vary ['veə.ri] [ver.i] - lišit se, různit se
to warrant ['wɒr.ənt] [wɔ:r-] - zaručit
trade-off ['treɪd.ɒf] [-ɑ:f] - kompromis
via [vaɪə] [vi:ə] - přes, prostřednictvím
wise (suffix) [waɪz] - pokud jde o, hovoříme-li o
write time [raɪt] [taim] - přístupová doba zápisu