

Storage Protection

KENICHI WADA
Hitachi Limited, Tokyo, Japan

5

Definition

Storage protection is a kind of data protection for data stored in a storage system. The stored data can be lost or becomes inaccessible due to, mainly, a failure in storage component hardware (such as a hard disk drive or controller), a disastrous event, an operator's mistake, or intentional alteration or erasure of the data. Storage protection provides the underlying foundation for high availability and disaster recovery.

Historical Background

In 1956, IBM shipped the first commercial storage that had a hard disk drive. To protect data from bit errors on disk platters, the hard disk drive commonly uses cyclic redundancy check (CRC) and an error-correcting code (ECC).

CRC and ECC cannot protect data from a whole disk failure in which an entire disk becomes inaccessible (for example, because of a disk head crash). The IBM 3990, which was shipped in the 1980s, had the replication functionality in which two identical copies of data were maintained on separate media. This approach protected data from this kind of failure. Replication functionality can be implemented in many other layers of the computer system. Most DBMS support database replication. Some file systems and Logical Volume Managers have file or volume replication functionality. Further, many storage systems and storage virtualization appliances support volume replication functionality.

RAID (Redundant Array of Inexpensive Disks) is another technology for protecting data from whole disk failure. D. Patterson et al. published a paper "A Case for Redundant Arrays of Inexpensive Disks (RAID)" in June 1988 at the SIGMOD conference [6]. This paper introduced a five level data protection scheme. The term RAID was adopted from this paper, but currently RAID is an acronym for Redundant Arrays of Independent Disks. It is noted that the patent covering RAID level 5 technology was issued in 1978 [5].

RAID level 1 is a kind of replication. RAID level 2 to 5 can reduce the capacity required to protect data against disk drive failure than replication, but it is limited to protect disk drive failure. Replication, on the other hand, can be used to protect databases, file systems and logical volume. Further replication can be used for disaster recovery, if data are replicated remotely.

Foundations

Hard disk drives commonly use Reed-Solomon code [7] to correct bit errors. Data in hard disk drives is usually stored in fixed length blocks. Controllers in hard disk drives calculate ECC for each block and record it associated with the original data. When data are read, the controller checks data integrity using ECC. CRC can be used with ECC for detecting bit errors and/or reducing the possibility of correction error.

Most DBMS support database replication with master/slave relation between the original and the replica. The master process updates and transfers it to the slave. This type of replication can provide high availability to the client of the DBMS in case of storage system failures as well as server failures. Another type of database replication is multi-master, which is mostly used to provide high performance parallel processing. Both types can be either synchronous or asynchronous replication. In synchronous replication, updates made in original are guaranteed in the replica, note there may be some delay in asynchronous replication.

Volume replication by storage system is also widely accepted as data protection. There are synchronous and asynchronous replications, the same as database replication. Asynchronous volume replication is often used for long distance remote replication. It may prevent performance degradation caused by replication delay, but could cause some data loss in case of recovery. Synchronous replication, on the other hand, may provide no data loss recovery, but may cause performance degradation due to replication delay. Volume replication is also used within a local datacenter for online backup. Backup servers use replica volume for backup during original volume is online. To support this, a storage system can pause update delegation from original to replica volume.

RAID (Redundant Array of Independent Disks) is a set of disks from one or more commonly accessible disk subsystems, combined with a body of control software, in which part of the physical storage capacity is used to store redundant information about user data stored on the remainder of the storage capacity. The term RAID refers to a group of storage schemes that divide and replicate data among multiple disks, to enhance the availability of data at desired cost and performance levels. A number of standard schemes have evolved which are referred to as levels. Originally, five RAID levels were introduced [6], but many more variations have evolved. Currently, there are several sublevels as well as many non-standard levels. There are trade-offs among RAID levels in terms of performance, cost and reliability.

Key Applications

- 115 Storage protection is essential to achieve business continuity and legal compliance with adequate performance, cost, and reliability.

(abridged)

120

Recommended Reading

1. ANSI. NFPA1600 Standard on Disaster/Emergency Management and Business Continuity Programs.
- 125 2. BSI. BS25999; Business Continuity Management.
3. Houghton A. Error Coding for Engineers. Kluwer Academic Publications, Hingham, MA, 2001.
4. Keeton K., Santos C., Beyer D., Chase J., and Wilkes J. Designing for disasters. In Proc. 3rd USENIX Conf. on File and Storage Technologies, 2004.
- 130 5. Ouchi N.K. (IBM Corporation). System for recovering data stored in failed memory unit. US Patent 4,092,732, 1978.
- 135 6. Patterson D., Gibson G., and Katz R. A case for redundant arrays of inexpensive disks (RAID). In 1988.
7. Sweeney P. Error Control Coding From Theory to Practice. Wiley, New York, 2002.
- 140 8. <http://www.sec.gov/>

Answer the following questions:

- 1) How can data loss occur?
- 2) What are CRC and ECC?
- 3) What does RAID stand for?
- 4) How many RAID levels were originally introduced?
- 5) What is Reed-Solomon code used for?
- 6) Besides disaster recovery, what is storage protection good for?
- 7) What is the basic difference between RAID 1 and the other RAID levels?

Match the following terms with their definitions:

- 1) synchronous replication
 - 2) asynchronous replication
 - 3) RAID level
 - 4) data integrity check
-
- a) in this approach, changes in the original can take some time to reflect in the replica
 - b) a way to determine whether data is corrupted
 - c) in this approach, changes in the original are immediately reflected in the replica
 - d) a specific strategy of distributing data over multiple disks

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

- 1) When data is read, the controller checks data integrity using CRC.
- 2) Replication can also be used to protect databases and file systems. Remote replication often employs the asynchronous volume replication approach.

Vocabulary

- appliance** [ə'plai.ənt s] 🗣️ - přístroj
- array** [ə'rei] 🗣️ - pole, sada
- compliance** [kəm'plai.ənt s] 🗣️ - shoda, dodržení
- degradation** [ˌdeg.rə'dei.ʃən] 🗣️ - pokles, zhoršení
- disastrous** [dɪ'zɑː.strəs] 🇺🇸 [-'zæs.trəs] 🗣️
- katastrofální, neblahý
- disk platter** [disk] 🗣️ ['plæt.ər] 🇺🇸 ['plæt.ə] 🗣️
- disková plotna
- erasure** [ɪ'reɪ.ʒər] 🇺🇸 [-ʒə] 🗣️ - smazání
- failure** ['feɪ.ljər] 🇺🇸 [ljə] 🗣️ - selhání
- inaccessible** [ˌɪn.ək'ses.i.bl] 🗣️ - nepřístupný,
nedostupný
- integrity** [ɪn'teg.rə.ti] 🇺🇸 [-tɪ] 🗣️ - celistvost,
neporušenost
- intentional** [ɪn'ten.ʃən.əl] 🗣️ - záměrný
- loss** [lɒs] 🇺🇸 [lɑːs] 🗣️ - ztráta
- recovery** [rɪ'kʌv.ər.i] 🇺🇸 [-ə-] 🗣️ - obnova
- redundancy** [rɪ'dʌn.dən.t.si] 🗣️ - nadbytečnost
- replica** ['rep.li.kə] 🗣️ - kopie, duplikát
- scheme** [ski:m] 🗣️ - schéma, soustava
- to evolve** [ɪ'vɒlv] 🇺🇸 [-'vɑːlv/] 🗣️ - vyvinout se
- volume** ['vɒl.ju:m] 🇺🇸 ['vɑːl-] 🗣️ - objem; svazek

Phrases

due to - kvůli; způsobený (čím)