

RAID (Redundant Array of Inexpensive Disks)

Illustration of RAID

RAID is a method of making magnetic disks fault tolerant. RAID uses an ECC (Error Correcting Code), known as Parity, to recover from single point of failure errors. That is, the failure of a single magnetic disk.

An Array (Stack) of Identical Magnetic Disks

RAID commonly uses from 3 to 7 magnetic disks. One of the magnetic disks is used for the parity information, and is not available for storing data. (In this example, the disks are numbered 1 through 5 with 5 being the parity disk)

[RAID spreads (stripes) the data and parity across all the disks to improve the speed of data retrieval, a capability not needed for almost all document management applications. For this reason, striping is not included in this description.]

For every bit on each of the identical disks, there is an identical bit on each of the other disks in the array. An example of a bit and the corresponding bits on the other disks is shown in the illustrations. The bits are linked with an arrow.

Bit Values

Each bit has a value of either a one or a zero.

Parity

With even parity, the total number of bits with a value of 1 is forced to be an even number by changing the value of the parity bit on the parity disk (disk number 5).

The total number of data bits (on disks 1 through 4) with a value of 1 is 3. To make the number of bits with a value of 1 an even number, the bit on the parity disk (disk 5) is changed to a 1. Now 4 bits have a value of 1.

Disk Failure

Disk 2 fails and is removed. To find the value of the bit identified by the arrow passing through disk 2, the corresponding bits on the other disks are read. 3 of these bits have a value of 1, so the value of the bit on the removed disk must have been a 1.

This same method can be used to determine the value of each of the other bits on the removed bit, completely reconstructing the data on that disk.

While the failed disk is removed, the data on the remaining disks must be read to reconstruct the data when a disk access is needed.

When a new disk is inserted, this method is used to rebuild the data on the new disk, without having to reload the data from a backup tape. When the new disk is fully rebuilt, the data can be read from the new disk.

Double Disk Failure

Parity does not work if more than one disk fails simultaneously. Parity only protects against a single point of failure.

ECC (Error Correcting Codes) on CDs, DVDs, and Magnetic Disks

CDs and DVDs have many bad bits (as many as one-half million per disk). CDs and DVDs use a much more elaborate, but analogous, error correcting code to correct these errors. This is why a newly copied CD or DVD is the best copy, because the ECC has just been used to generate a perfect copy of the data and ECC to write to the new CD or DVD. As the CD or DVD deteriorates over time, the number of errors (bad bits) increases. When the number of bad bits exceeds the number that can be corrected by the ECC (two bad bits (on two disks) are enough to overwhelm the parity ECC described above), a catastrophic failure occurs and no bits can be read from the disk. This is why CDs and DVDs must be copied before the number of bad bits build up to a level that exceeds what the ECC can protect from. ANSI / AIIM (Association for Information and Image Management / American National Standards Institute)MS59 1996 is a standard that provides for reading the level of bad bits on a disk. MS 59 complies with the ANSI X3.131 media error hardware interface standard. [Individual magnetic disks also have a robust ECC that is used to correct data bits before the bits are given to the RAID parity ECC.] [The Internet also uses codes to detect and correct errors that occur during data transmission.]

Note to Readers

Updates and More Detailed Descriptions

When using the information in this article, please check the website <http://www.ArchiveBuilders.com> for updates. The version number of this article is just before the page number below. The website also has articles that provide more details on some of the terms and concepts in this article.

Comments

Please let us know how you like this paper, or if you had any questions. What would you like to see in the future? Also, please let us know where you saw this paper. For more, and the most recent version of this article, please visit our web site at <http://www.ArchiveBuilders.com>

Please send your comments via email to SteveGilheany@ArchiveBuilders.com. Tel: +1 (310) 937-7000 Fax: +1 (310) 937-7001. Also, please let us know where you saw this article.

Acknowledgements

Reprinted from *Archive Planning*, Volume 4, number 3, 2000, Archive Builders' analysis newsletter for document management.

See <http://www.ArchiveBuilders.com>

All trademarks are the property of their respective holders.

Note to Editors

Paper 22020v009

We will continue to update these articles as we get comments. Please contact us for the most current version before you publish. Also, please request permission to publish the article. Permission will be given freely for most purposes.

Steve Gilheany
Archive Builders
1209 Manhattan Ave., C-14
Manhattan Beach, CA 90266
Tel: +1 (310) 937-7000 Fax: +1 (310) 937-7001
SteveGilheany@ArchiveBuilders.com

Dividing this Article into Parts for Serialization

If you decide to divide this article into parts please print at least the updates, comments, and acknowledgements sections in each of the parts along with:

'by SteveGilheany@ArchiveBuilders.com'

Bio

Steve Gilheany, BA in Computer Science, MBA, MLS Specialization in Information Science, CDIA (Certified Document Imaging System Architect), AIIM Master (MIT), and AIIM Laureate (LIT), of Information Technologies, CRM (Certified Records Manager, ARMA) has eighteen years experience in document imaging and is a Sr. Systems Engineer at Archive Builders.

Author

Steve Gilheany is a Sr. Systems Engineer at Archive Builders. He has worked in digital document management and document imaging for twenty years.

His experience in the application of document management and document imaging in industry includes: aerospace, banking, manufacturing, natural resources, petroleum refining, transportation, energy, federal, state, and local government, civil engineering, utilities, entertainment, commercial records centers, archives, non-profit development, education, and administrative, engineering, production, legal, and medical records management. At the same time, he has worked in product management for hypertext, for windows based user interface systems, for computer displays, for engineering drawing, letter size, microform, and color scanning, and for xerographic, photographic, newspaper, engineering drawing, and color printing.

In addition, he has nine years of experience in data center operations and database and computer communications systems design, programming, testing, and software configuration management. He has an MLS Specialization in Information Science and an MBA with a concentration in Computer and Information Systems from UCLA, a California Adult Education teaching credential, and a BA in Computer Science from the University of Wisconsin at Madison. His industry certifications include: the CDIA (Certified Document Imaging System Architect) and the AIIM Master (MIT), and AIIM Laureate (LIT), of Information Technologies (from AIIM International, the Association of Information and Image Management, www.AIIM.org), and the CRM (Certified Records Manager) (from the ICRM, the Institute of Certified Records Managers, an affiliate of ARMA International, the Association of Records Managers and Administrators, www.ARMA.org).

Contact:

SteveGilheany@ArchiveBuilders.com
Tel: +1 (310) 937-7000 Fax: +1 (310) 937-7001

RAID (Redundant Array of Inexpensive Disks)

