Scientific Journal Impact Factor (SJIF): 3.518
International Journal of Modern Trends in Engineering and Research
ISSN (ONLINE) : 2349 - 9745
ISSN (PRINT) : 2393 - 8161

RAID: The Innovative Data Storage Manager

Amit Tyagi
IIMT College of Engineering, Greater Noida, UP, India

Abstract-RAID is a technology that is used to increase the performance and/or reliability of data storage. The abbreviation stands for *Redundant Array of Inexpensive Disks*. A RAID system consists of two or more disks working in parallel. These disks can be hard discs but there is a trend to also use the technology for solid state drives. There are different RAID levels, each optimized for a specific situation. These are not standardized by an industry group or standardisation committee. This explains why companies sometimes come up with their own unique numbers and implementations.

In the late 1980s, researchers at the University of California at Berkeley were looking for ways of combining disks into arrays with desirable combinations of affordability, data availability and I/O performance.

In 1988, a landmark paper entitled A Case for Redundant Arrays of Inexpensive Disks was published. The paper demonstrated that arrays of low-cost personal computer disks could be effective substitutes for the high capacity, high performance disks then used with data centre computers. Five disk array models, called RAID Levels, were described.

I. Introduction

RAID is a mature technology that speeds up data access while at thesame time protecting your data from hard disk failure. RAID is quicklybecoming a necessary component in every network since data lossand downtime can prove both fatal and financially destructive. Mostnetworks are designed to provide instant access to massive amountsof data. More and more employees have to access customer and otherdatabases. Intranets and corporate Web sites provide access to hugedatabases online.

The software to perform the RAID-functionality and control the hard disks can either be located on a separate controller card (a hardware RAID controller) or it can simply be a driver. Some versions of Windows, such as Windows Server 2003, as well as Mac OS X include software RAID functionality. Hardware RAID controllers cost more than pure software but they also offer better performance.

RAID-systems can be based with a number of interfaces, including SCSI, IDE, SATA or FC (fibre channel.) There are systems that use SATA disks internally but that have a FireWire or SCSI-interface for the host system.

Sometimes disks in a RAID system are defined as JBOD, which stands for *'Just a Bunch Of Disks'*. This means that those disks do not use a specific RAID level and acts as stand-alone disks. This is often done for drives that contain swap files or spooling data.

II. Raid Components and Concepts

First, let us define Logical Arraysas a split or combination of Physical Arrays, which in turn are one or more Physical Drives that are simply the individual hard disks that comprise these arrays. Logical Drives are then made of one or more Logical Arrays.

Mirroring refers to complete redundancy of data on identical disks. The data that is being written on one Logical Array is completely duplicated on a similar array thereby providing 100% data redundancy. The cost associated with mirroring is that the amount of available storage is reduced by 50%; writes are slightly slower albeit reads are faster in some situations.

Striping refers to a technique that allows Physical Drives in a Logical Array to be used in parallel in order to gain in performance. In this technique, data is broken down in Byte or Block levels or stripes, where every Byte or Block is written to a separate disk in the array. Byte level can at times be a 512-byte sector, while Block size can be selected from variety of choices. The gain in performance is similar between Reads and Writes.

In some RAID levels, striping is combined with a technique called Parity to enhance fault tolerance. Parity, similar to parity in memory, is simply adding a Block (Byte) of calculated parity data to several Blocks (Bytes) in such a way that any one of the Blocks (Bytes) can be reconstructed in case of loss, from the remainder of the Blocks (Bytes) and the parity Block (Byte). While Parity gains from performance of striping, its disadvantages are more complexity and loss of some disk space, which is taken up by parity information. There are many ways to combine RAID techniques.

Some standardized combinations are referred to as RAID Levels, even though ‘Level’ in this context does not denote any hierarchy or advantage. Levels are independent and different. Some RAID levels combine multiple other levels to achieve certain aims.

The RAID Advisory Board (RAB) has been active since 1992 in education and standardization of RAID technology. See **http://www.raid-advisory.com/.** Techniques discussed above are used in different levels. Mirroring is used in levels 1, 0+1, 10 (1+0). Striping without parity is used in level 0, 0+1, and 10. Striping with Block level parity is used in level 5 and 6. While the minimum number of drives required at each level are noted, there is no inherent maximum to number of drives in arrays other than the one imposed by controllers.

III. RAID Levels

RAID level 0 – Striping

In a RAID 0 system data are split up in blocks that get written across all the drives in the array. By using multiple disks (at least 2) at the same time, this offers superior I/O performance. This performance can be enhanced further by using multiple controllers, ideally one controller per disk.

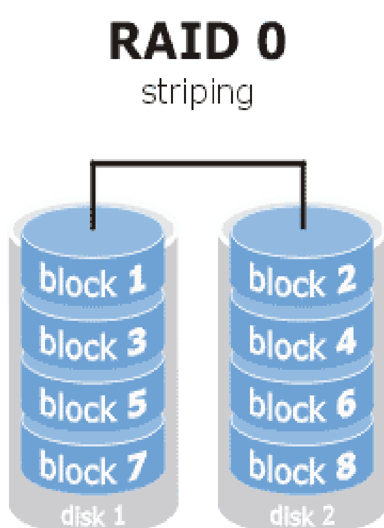

Advantages

- RAID 0 offers great performance, both in read and writes operations. There is no overhead caused by parity controls.
- All storage capacity is used, there is no disk overhead.

- The technology is easy to implement.

Disadvantages

RAID 0 is not fault-tolerant. If one disk fails, all data in the RAID 0 array are lost. It should not be used on mission-critical systems.

Ideal use

RAID 0 is ideal for non-critical storage of data that have to be read/written at a high speed, such as on a Photoshop image retouching station

RAID level 1 – Mirroring

Data are stored twice by writing them to both the data disk (or set of data disks) and a mirror disk (or set of disks) . If a disk fails, the controller uses either the data drive or the mirror drive for data recovery and continues operation. You need at least 2 disks for a RAID 1 array.

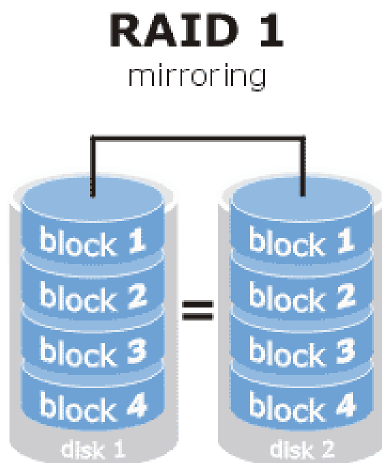

RAID 1 systems are often combined with RAID 0 to improve performance. Such a system is sometimes referred to by the combined number: a RAID 10 system.

Advantages

- RAID 1 offers excellent read speed and a write-speed that is comparable to that of a single disk.
- In case a disk fails, data do not have to be rebuilding, they just have to be copied to the replacement disk.
- RAID 1 is a very simple technology.

Disadvantages

- The main disadvantage is that the effective storage capacity is only half of the total disk capacity because all data get written twice.
- Software RAID 1 solution do not always allow a hot swap of a failed disk (meaning it cannot be replaced while the server keeps running). Ideally a hardware controller is use

Ideal use

RAID-1 is ideal for mission critical storage, for instance for accounting systems. It is also suitable for small servers in which only two disks will be used.

RAID level 3

On RAID 3 systems, data blocks are subdivided (striped) and written in parallel on two or more drives. An additional drive stores parity information. You need at least 3 disks for a RAID 3 array.

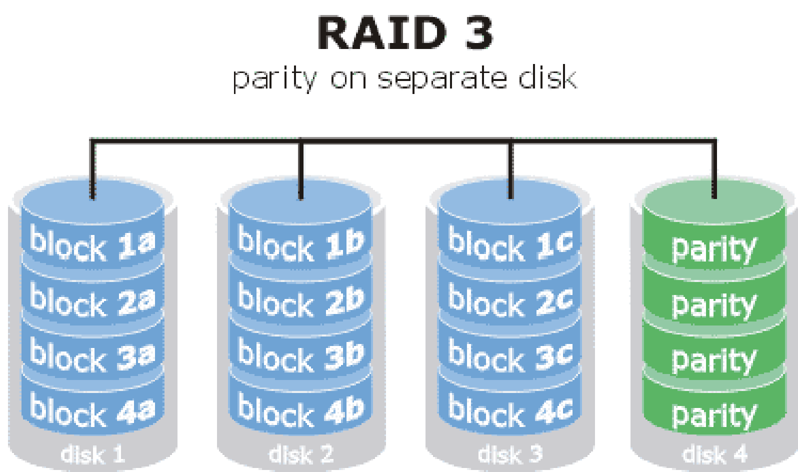

Since parity is used, a RAID 3 stripe set can withstand a single disk failure without losing data or access to data.

Advantages

- RAID-3 provides high throughput (both read and write) for large data transfers.
- Disk failures do not significantly slow down throughput.

Disadvantages

- This technology is fairly complex and too resource intensive to be done in software.
- Performance is slower for random, small I/O operations.

Ideal use

RAID 3 is not that common in prepress.

RAID level 5

RAID 5 is the most common secure RAID level. It is similar to RAID-3 except that data are transferred to disks by independent read and write operations (not in parallel). The data chunks that are written are also larger. Instead of a dedicated parity disk, parity information is spread across all the drives. You need at least 3 disks for a RAID 5 array.
A RAID 5 array can withstand a single disk failure without losing data or access to data. Although RAID 5 can be achieved in software, a hardware controller is recommended. Often extra cache memory is used on these controllers to improve the write performance.

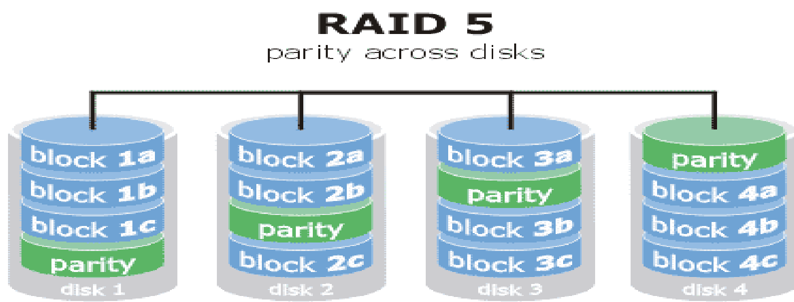

Advantages

Read data transactions are very fast while write data transaction are somewhat slower (due to the parity that has to be calculated).

Disadvantages

- Disk failures have an effect on throughput, although this is still acceptable.
- Like RAID 3, this is complex technology.

Ideal use

RAID 5 is a good all-round system that combines efficient storage with excellent security and decent performance. It is ideal for file and application servers.

RAID level 10 – Combining RAID 0 & RAID 1

RAID 10 combines the advantages (and disadvantages) of RAID 0 and RAID 1 in one single system. It provides security by mirroring all data on a secondary set of disks (disk 3 and 4 in the drawing below) while using striping across each set of disks to speed up data transfers.

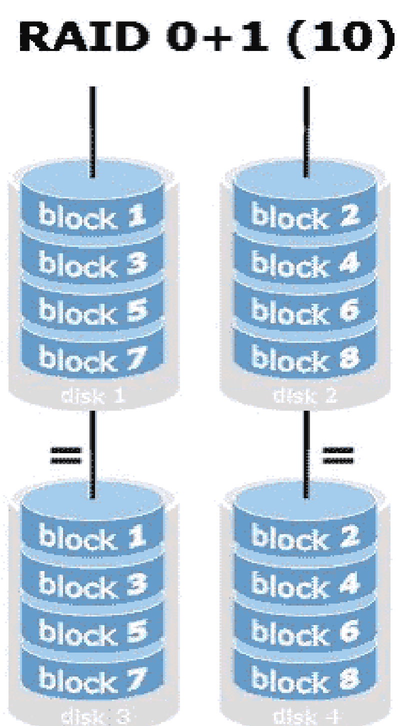

What about RAID levels 2, 4, 6 and 7?

These levels do exist but are not that common, at least not in prepress environments. This is just a simple introduction to RAID-system.

RAID is no substitute for back-up!

All RAID levels except RAID 0 offer protection from a single drive failure. A RAID 6 system even survives 2 disks dying simultaneously. For complete security you do still need to back-up the data from a RAID system.

- That back-up will come in handy if all drives fail simultaneously because of a power spike.
- It is a safeguard if the storage system gets stolen.
- Back-ups can be kept off-site at a different location. This can come in handy if a natural disaster or fire destroys your workplace.
- The most important reason to back-up multiple generations of data is user error. If someone accidentally deletes some important data and this goes unnoticed for several hours, days or weeks, a good set of back-ups ensure you can still retrieve those files.

IV. Summary: Why Is RAID Important?

As the storage industry becomes increasingly independent of the computer system industry, storage alternatives are becoming more complex. System administrators, as well as managers who make storage purchase and configuration decisions need to understand on-line storage alternatives. Awareness of what RAID can and cannot do for them helps managers make informed decisions about on-line storage alternatives.

Users of networked personal computers may also be concerned about the quality of the storage service provided by their data servers. The material in this paper can help make the personal computer user aware of the significance of available storage alternatives.

Moreover, the first desktop RAID systems have already appeared on themarket. As disk size and cost continue to decline, widespread use of RAID on the desktop is only a matter of time.

The purpose of this paper has been to provide the background to help users formulate storage subsystem strategies, particularly with respect to RAID subsystems. The market's view has progressed from RAID as "add-on extra," through RAID as "necessity in mission-critical applications," to "RAID unless there's a good reason not to." In 1997, according to one analyst, about 80% of the disk subsystems shipped was RAID-capable. The analyst predicted RAID ubiquity by the year 2000. It appears that even today, the average server disk is part of a RAID array; non-arrayed disks are the exception rather than the rule.

REFERENCES

[1] "Raid for Enterprises Computing": Veritas Technology's White Paper.
[2] "Raid Explained": By Jerry Scott.
[3] "Raid Levels 0,1,2,3": By Google Inc.