RAID - Redundant Array of Independent Disks

History

Most PC systems have only one hard drive or disk. If a hard disk fails in single disk systems (or multiple disk systems where each drive is ‘standalone’) then the data on that disk is lost. If regular backups of the data have been made then the data loss may be minimal, but restoring data from backups can be a difficult and time-consuming process. RAID technology incorporates ‘redundancy’ – the duplication of critical components in order to improve the overall reliability of the system – and therefore removes the ‘single point of failure’ inherent in non-RAID systems. The use of multiple disks can also improve system performance by sharing read and write operations across disks, thus making optimum use of system resources.

How RAID works

RAID provides redundancy by storing the same data in different places. Identical sets of data are stored on multiple disks, ensuring there is always a back-up copy should one disk fail. A RAID appears to the operating system to be a single hard disk (or logical volume). The RAID controller manages the multiple disks so that the operating system works without knowing anything about the number or condition of disks ‘behind’ the RAID controller. This means that the operating system continues to work even when a disk fails; the RAID controller ensures that the damaged or missing information is supplied using the backup or ‘redundant’ data.

There are various types of RAID configuration, each designed to address a different range of applications. There are applications that require a high level of performance (i.e. very fast read/write operations), such as video editing, and others that require a high level of redundancy or security, such as an operating system or an application which stores valuable data (e.g. family photos, music). There is even one configuration that provides no redundancy at all – RAID-0 – which is designed purely to improve performance.

There are two basic RAID techniques: ‘mirroring’ and ‘striping’.

Mirroring

Mirroring, or ‘disk mirroring’, simply refers to the use of at least two disks to store duplicate data. The contents of one disk are mirrored by the other disk, ensuring that a single disk failure will not cause data loss as there will always be a ‘good’ or redundant copy on the other drive. Once the failed drive is replaced, the RAID controller will rebuild the array using the information on the good disk.

Compare the diagrams below showing an overview of a track being written on a system without RAID (normally referred to as ‘IDE’) (fig. 1) and a system with RAID-1 (fig. 2).

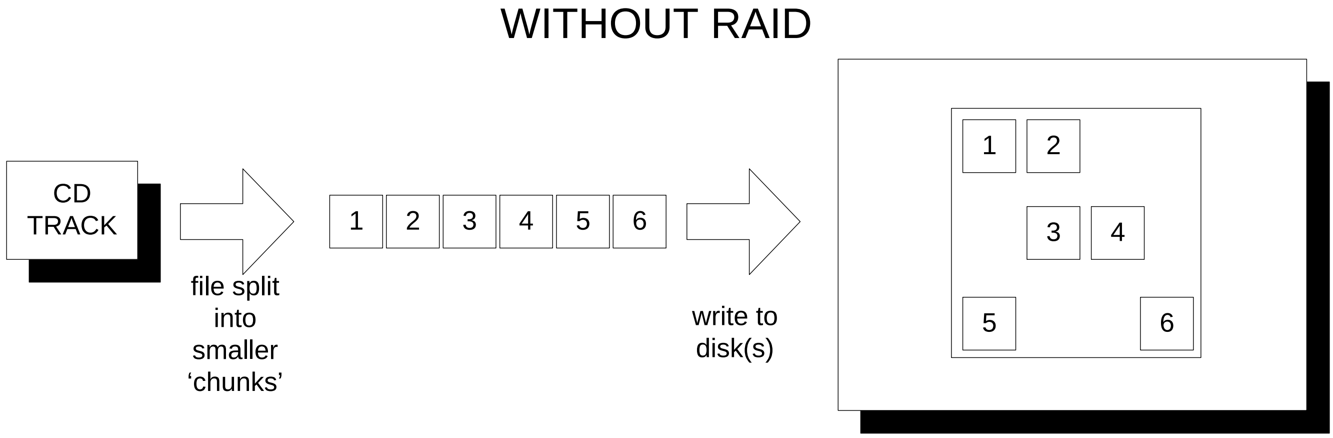

Figure 1 – overview of ‘normal’ CD track write

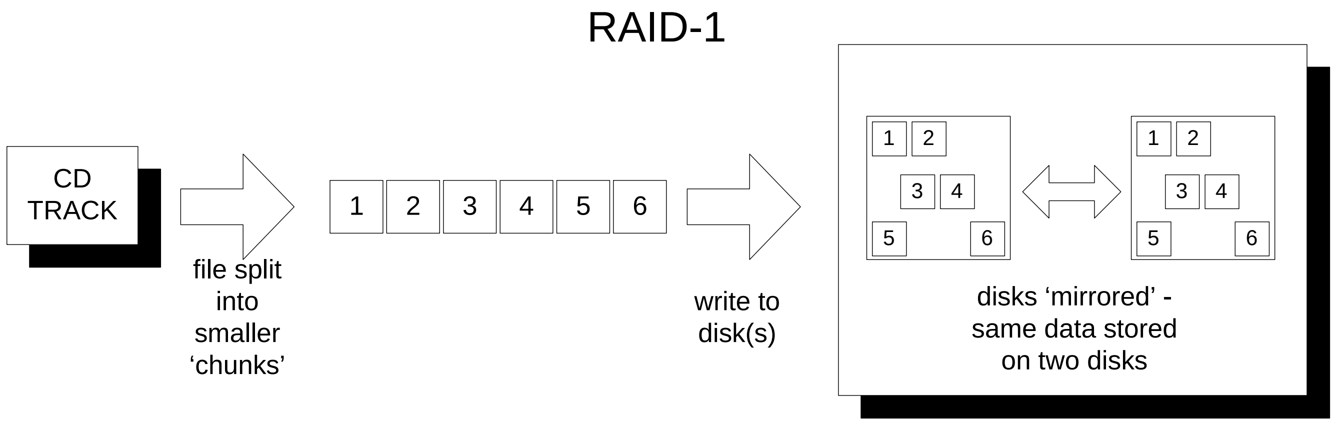

Figure 2 – overview of RAID-1 CD track write

RAID-1 uses 50% of the disk capacity in order to provide redundancy, reducing signifcantly the overall capacity of the system. There is also an effect on write performance as two 'copies' must be created every time data is written. However, there is a performance improvement in read operations as both disks can be read from simultaneously: each disk retrieves a different part of the data and the RAID controller forms the complete file.

Striping
Striping refers to the technique of dividing the data to be written into equally sized chunks or 'stripes' that are then written simultaneously to multiple disks. Striping provides a significant performance improvement as each disk can be reading or writing at the same time; rather than waiting for one disk to read or write one large file, the system is waiting only long enough for multiple disks to read or write each stripe. Striping does not provide any redundancy by itself; configurations that only use striping, such as RAID-0, do not therefore improve data integrity or reliability.

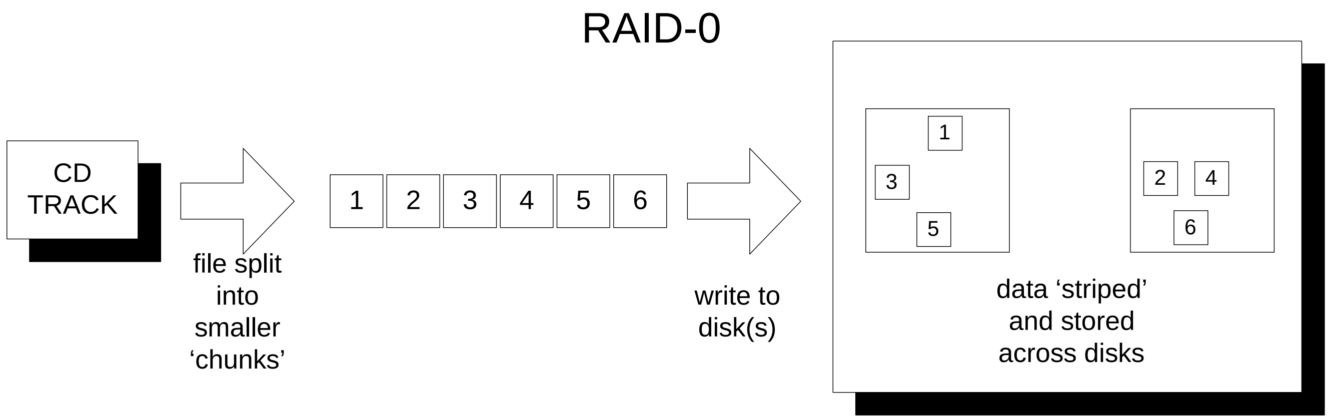

Figure 3 – overview of RAID-0 CD track write

How the KIVOR TUNBOKS/INDEX use RAID

The KIVOR TUNBOKS/INDEX uses a RAID-1 configuration to store the operating system (the basic software required to operate, including the RAID controller) and RAID-5 to store the audio data:

RAID-1 – Mirrored Set
RAID-1 uses two disks to store duplicate operating system data. The KIVOR TUNBOKS/INDEX is still able to boot up and operate normally even if one of the disks fails. RAID-1 is used for operating system data for the following reasons: it provides the highest level of redundancy and therefore data integrity; the 50% storage overhead is not a big issue because the operating system does not require a large amount of disk space; and finally, the slight drop in write performance is not a problem as write operations are relatively few.

RAID-5 – Striped Set with Distributed Parity
RAID-5 stripes the data across three disks and uses the fourth disk to store parity information. Parity information is not an exact mirrored copy of the data but is created from it by the RAID controller and can be used to exactly re-create the lost data in the event of a failure. The actual disk that stores the parity information changes with each write operation; the use of 'distributed' parity ensures that read and write operations can be overlapped in a way that is not possible with a dedicated parity disk (used by RAID-4).

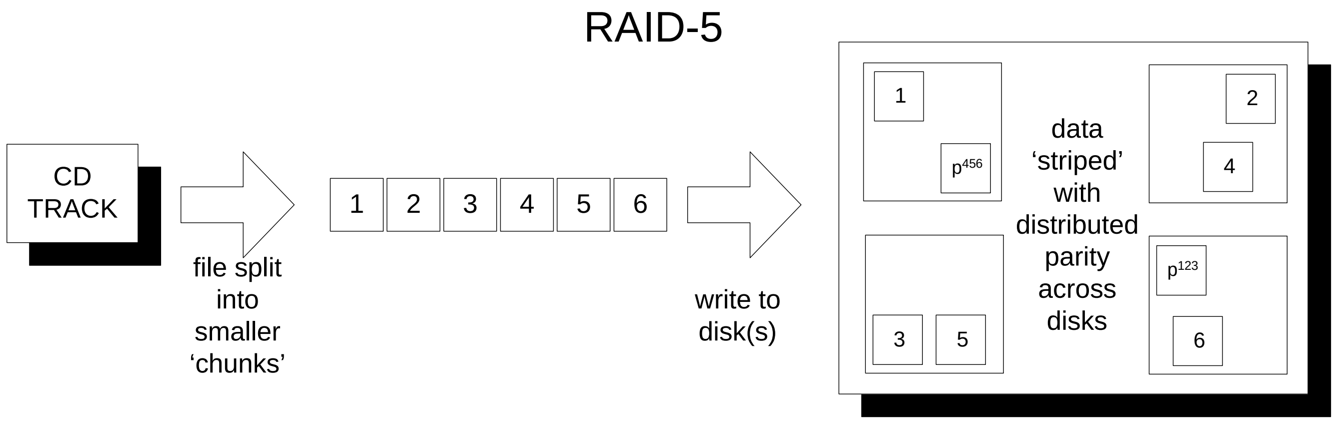

Figure 4 – overview of RAID-5 CD track write

The RAID controller creates the parity information by performing a simple XOR operation on each individual data bit. If any of the bits are lost or damaged, the procedure can be reversed to restore the data from the parity information and the other data bits (this is done ‘on the fly’ by the RAID controller until the array is re-built). The parity information is equal in size to the ‘stripe size’ (the size of the ‘chunk’ the data is broken into). With a stripe size of 64K, a 192K file is divided into three 64K stripes with each disk receiving one stripe and the fourth disk receiving the 64K parity file. There is a 25% storage overhead when using RAID-5 (compared to 50% with RAID-1) because of the space the parity information takes up; therefore, the total storage space of the system will be 75% of the sum of the disk capacities.

RAID-5 provides a good balance between security and performance: the use of striping brings advantages in read/write performance; the use of parity information allows more of the disk to be used to store audio data without compromising on the level of redundancy.

Summary

- RAID configurations (excluding RAID-0) improve data integrity and fault tolerance
- RAID can improve performance in storage-intensive applications
- RAID-1 uses 50% of total storage space in order to provide redundancy
- RAID-5 uses 25% of total storage space in order to provide redundancy
- RAID-5 can continue to operate even with a single drive failure