

# Understanding RAID?

## Overview

This document is aimed at the CEO/COO, Managing Director and Directors and Senior Managers level. It is not aimed at the IT Manager. The document aims to provide you with some basic knowledge about this vital subject and to help you understand what your choices are for ensuring that all your companies data is stored in a safe way. It explains where the problems can be and what you need to do about finding a solution.

Before we define what RAID is, it is worth understanding how a standard computer system stores data. In this document all references computers systems refer to servers. For the purposes of this document the operating system that these servers run is not relevant.

Servers are used by companies to store data on, that data can be your companies accounts system or an Enterprise Resource Planning System (ERP), your e-mail system or just simply files, such as spreadsheets and Word documents held in a central, series of directories or at its most simple a share directory visible to all people on your network.

So we now have 1 or more servers holding large quantities of your companies most important data. When you are running your accounts or e-mail on a server it is of vital importance that this data is not lost. Computers store data on hard disks. What is a hard disk? A hard disk (commonly known as a HDD (hard disk drive) or hard drive (HD) is a non-volatile storage device which stores digitally encoded data on rapidly rotating platters with magnetic surfaces. See below for a picture of a hard disk with its cover removed, this is a typical server hard disk.



Because hard disks are electro mechanical devices there is always the chance of failure, in the 1<sup>st</sup> year of use the failure rate will typically be in the 4% to 6% range there after it rises to around 8% to 10% per year.

So even with a modest number of hard disks most companies are likely to see one hard disk fail a year.

## Why do I need raid on my server?

---

When a hard disk fails the data contained on that hard disk is now inaccessible. Yes data can be recovered but that takes time, usually several days, and can cost more than the cost of the actual disk itself.

So we need a method of protecting our data that is stored on the hard disks. This is what RAID sets out to achieve.

### **RAID Definition** (definition supplied by Wikipedia.com)

The acronym RAID (originally **R**edundant **A**rray of **I**nexpensive **D**rives (or disks), refers to a data storage scheme using multiple hard drives to share or replicate data among the drives. Depending on the configuration of the RAID (typically referred to as the RAID level), the benefit of RAID is to increase data integrity, fault-tolerance, throughput and/or capacity, compared with single drives. In its original implementations, its key advantage was the ability to combine multiple low-cost devices using older technology into an array that offered greater capacity, reliability, speed, or a combination of these things, than was affordable available in a single device using the newest technology.

Fundamentally, RAID combines multiple hard disks into a single logical unit. There are two ways this can be done: in hardware and in software. Hardware combines the drives into a logical unit in dedicated hardware which then presents the drives as a single drive to the operating system. Software does this within the operating system and presents the drives as a single drive to the users of the system.

RAID is typically used on servers but can be used on workstations, such as CAD stations and Graphic workstations. This is especially true in storage-intensive computers such as video and audio editing.

In a multi-user system, better performance requires establishing a stripe wide enough to hold the typical or maximum size record. This allows overlapped disk I/O across drives.

A quick summary of the most commonly used RAID levels (there are other levels, however the ones below are the most common in use today):

**RAID 0: Striped Set (2 disk minimum) without parity:** provides improved performance and additional storage but no fault tolerance from disk errors or disk failure. Any disk failure destroys the array, which becomes more likely with more disks in the array.

**RAID 1: Mirrored Set (2 disks minimum) without parity:** provides fault tolerance from disk errors and single disk failure. Increased read performance occurs when using a multi-threaded operating system that supports split seeks, very small performance reduction when writing. Array continues to operate with one failed drive.

**RAID 3 and RAID 4: Striped Set (3 disk minimum) with Dedicated Parity:** Provides improved performance and fault tolerance similar to RAID 5, but with a dedicated parity disk rather than rotated parity stripes. The single disk is a bottle-neck for writing since every write requires

## Why do I need raid on my server?

---

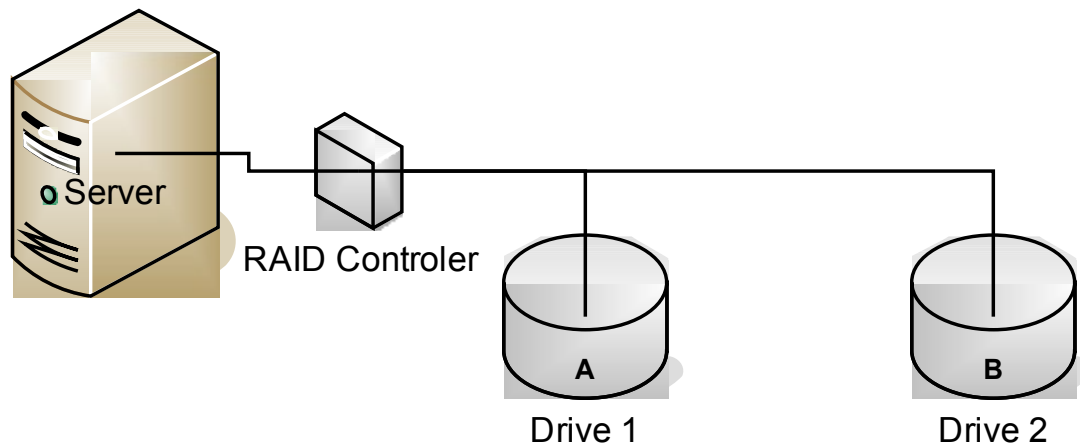
updating the parity data. One minor benefit is the dedicated parity disk allows the parity drive to fail and operation will continue without parity or performance penalty.

**RAID 5: Striped Set (3 disk minimum) with Distributed Parity:** Distributed parity requires all but one drive to be present to operate; drive failure requires replacement, but the array is not destroyed by a single drive failure. Upon drive failure, any subsequent reads can be calculated from the distributed parity such that the drive failure is masked from the end user. The array will have data loss in the event of a second drive failure and is vulnerable until the data that was on the failed drive is rebuilt onto a replacement drive.

**RAID 6: Striped Set (4 disk minimum) with Dual Distributed Parity:** Provides fault tolerance from two drive failures; array continues to operate with up to two failed drives. This makes larger RAID groups more practical. This is becoming a popular choice for SATA drives as they approach 1 Terabyte in size. This is because the single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt. The larger the drive, the longer the rebuild will take. With dual parity, it gives the array time to rebuild onto a large drive with the ability to sustain another drive failure.

Now that we understand the basics of RAID, we will look at the 3 most common systems of RAID that are implemented in server systems today. These are RAID: 0, 1, and 5. Most servers today ships with RAID 0 and 1 support as standard. This means that it is already built in to the server hardware and only requires you to add the disk and say if you want RAID 0 or 1. However in order to add RAID 5 support to a server you need to purchase an additional item of hardware called a RAID 5 card.

### RAID 0



Requires at least 2 disks, any number of additional disk can be added.

#### Characteristics & Advantages

RAID 0 implements a striped disk array, the data is broken down into blocks and each block is written to a separate disk drive

#### Disadvantages


Not a "True" RAID because it is NOT fault-tolerant

The failure of just one drive will result in all data in an array being lost



# Why do I need raid on my server?

## RAID 5



Each entire data block is written on a data disk; parity for blocks in the same rank is generated on Writes, recorded in a distributed location and checked on Reads. Distributed parity requires all but one drive to be present to operate; drive failure requires replacement, but the array is not destroyed by a single drive failure. Upon drive failure, any subsequent reads can be calculated from the distributed parity such that the drive failure is masked from the end user. The array will have data loss in the event of a second drive failure and is vulnerable until the data that was on the failed drive is rebuilt onto a replacement drive.

RAID Level 5 requires a minimum of 3 drives to implement

### Characteristics & Advantages

Highest Read data transaction rate

Medium Write data transaction rate

Low ratio of ECC (Parity) disks to data disks means high efficiency

Good aggregate transfer rate

### Disadvantages

Disk failure has a medium impact on throughput

Most complex controller design

Difficult to rebuild in the event of a disk failure (as compared to RAID level 1)

Individual block data transfer rate same as single disk

## Conclusion

RAID 0 is very good for speed as it writes and reads data from all disks in the array, but in the event of a disk failure every thing is lost. This is the cheapest to implement but offer no redundancy, just high speed.

RAID 1 mirrors the data so that the data is effectively twice, once to each disk. This means that in the event that 1 disk fails the other disk is an exact copy and you can carry on using the system. This offers redundancy and speed but uses double the number of disk drives

RAID 5 This offers very good redundancy as well as reasonable speed. No drives are wasted as they are aggregated together and presented to the system as 1 large drive. Most complicated to set-up. In the event of a failure of the RAID card, all the data can be lost.

If you want the most security for your data then you should use RAID 5, however you must also have some form of backup so that you can make a copy of your data in the event that there is a total failure of the RAID system or the Server. Please see the Comware White Paper concerning Backup.

If data security is of critical importance to you you might want to consider using clustered servers. This is the subject of an additional White Paper.