

## What is RAID

**Submitted By:** Jason Mullett **Date:** 10/19/2003

### Redundant Array of Independent Disks (RAID)

RAID is the term for an array of two or more disk employed in combination for performance and redundancy (fault tolerance). The term was coined in 1988 in a paper describing array configuration and application by researchers and authors Patterson, Gibson and Katz of the University of California at Berkley. A single hard drive in a computer is the weakest link in the chain and is more prone to failure than any other part of a computer. RAID was developed in response to this problem. RAID, depending on the configuration level used will improve read and write performance and provide redundancy.

#### RAID Terms:

**Redundancy** Fault tolerance, protects data from hard disk failure.

**Mirroring** The same data is written to two hard drives, providing fault tolerance but decreasing write performance. (i.e. it would not be a good policy to place a page file on a mirrored volume)

**Parity Bit** An odd or even data bit that defines the value of the missing data bit should a hard drive drop out of an array that uses parity for redundancy. If the added value of a data block is even than a "0" is generated as the parity bit if the added value is odd the parity bit is a "1". Depending on the RAID level the parity will either be on one disk or be spread among all the disks. Either way it is 1/5th or 20% of the space when you are utilizing five disks. Parity is 1/4th or 25% of the space when utilizing four disks, and 1/3rd or 33% when utilizing three disks. The diagram below better describes how parity works.

	Disk 1	Disk 2	Disk 3	Disk 4	Disk 5
<b>Data block 1</b>	parity=0	1	0	0	1
<b>Data block 2</b>	0	parity=1	1	1	1
<b>Data block 3</b>	1	1	parity=0	0	0
<b>Data block 4</b>	0	1	1	parity=1	1
<b>Data block 5</b>	0	0	1	0	parity=1

**Striping** Striping breaks up a file and writes it to multiple disks concurrently. Striping can be done at either a block level (the size is determined by the administrator), byte level, or bit level.

#### RAID Configuration Levels:

**RAID0** Data striping Array

**RAID1** Mirrored Disk Array

**RAID2** Parallel Array, Hamming Code


**RAID3** Parallel Array with Parity

**RAID4** Independent Actuators with a dedicated Parity Drive

**RAID5** Independent Actuators with parity spread across all drives

**RAID0 Defined**


**RAID 0**



RAID0 is a striped disk array without parity. The data is broken into blocks and is written across two or more disk drives. RAID0 greatly improves disk performance, but provides no fault tolerance and thus is not a true form of RAID. RAID0 is the only RAID configuration level that does not provide redundancy.

**RAID1 Defined**


**RAID 1**



Mirrored volumes or duplicated disks. A RAID1 volume is two disks with duplicate data. The data is written to two separate hard disks. RAID1 or mirrored volumes provide redundancy, but cause a decrease in performance because the data is written twice. Read performance does not increase or decrease, but is equal to that of one disk. (i.e. it would not be a good policy to place a page file on a mirrored or RAID1 volume)

**RAID2 Defined**

**RAID 2**




A0 to A3=Word A; B0 to B3 = Word B; C0 to C3=Word C; D0 to D3 = Word D  
 ECC/Ax to Az=Word A ECC; ECC/Bx to Bz = Word B ECC;  
 ECC/Cx to Cz=Word C ECC; ECC/Dx to Dz = Word D ECC

Each bit of data word is written to a data disk drive (4 in this example: 0 to 3). Each data word has its Hamming Code ECC word recorded on the ECC disks. On Read, the ECC code verifies correct data or corrects single disk errors. RAID2 provides "on the fly error correction" and makes extremely high data transfer rates possible. RAID2 is not commercially available.

**RAID3 Defined**


**RAID 3**



On RAID3 the data block is striped or written onto the data disks. The parity bit is generated during the write and is written to the parity disk and checked on reads. RAID3 provides significantly improved read and write data transfer rates and provides redundancy. RAID3 requires a minimum of three drives to implement.

**RAID4 Defined**


**RAID 4**



The entire data block is written onto a data disk. The parity is generated on the write and written to the parity disk. RAID4 has high read transaction rate, but a poor write transaction rate. RAID4 requires a minimum of three drives to implement.

**RAID5 Defined**

**RAID 5**



RAID5 is the most popular RAID of choice. RAID5 stripes the data blocks across three or more disks and records the parity bit in distributed locations. RAID5 provides a high read transaction rate and a medium write transaction rate. RAID5 also provides redundancy.