

Making Sense of RAID

By: Adam Chee W.S

Note: This article is also published at [TechTutorials.com](#)

Whether you are a student preparing for your Microsoft certifications, a System Administrator maintaining a production server or a System Engineer developing a new solution, an understanding of RAID technology is nothing but beneficial.

As a System Administrator, an understanding of RAID technology implemented in the servers you maintain will aid you in your troubleshooting processes. As a System Engineer, it will help you to develop a more efficient and robust system and as a student, it will help a little in preventing you from becoming a paper MCSA/MCSE.

In today's context where the lost data or information would directly or indirectly affect your profit margin or competitive edge, RAID has become the norm in most mission-critical implementations. RAID or Redundant Array of Independent Disks is really just an array of multiple hard disks 'grouped' together to achieve benefits not possible with a single hard disk. To the host computer, the RAID array appears as a single storage unit and this 'grouping' of hard disks can be performed either via hardware or software.

Hardware versus Software

As mentioned earlier, there are two ways to achieve RAID implementations, either by hardware or via software, for software based RAID, I'll be touching only on Windows 2000 supported software RAID.

Hardware RAID

- Hardware based RAID is managed independently from the host.
- Performance of the RAID depends on the controller device's limitation (E.g. SCSI or Fiber Channel) but it would still be faster I/O than software RAID
- True hardware multi-tasking and fault tolerant as the host CPU and the RAID can run in parallel
- More expensive than software RAID as additional hardware is needed.

Software RAID

- Software RAID is application-based and requires the host's memory as well as and CPU cycles, thus degrading the overall server performance.
- Performance of the RAID depends on the host's CPU performance and load.
- If the RAID application fails, the software RAID might not boot up.
- Usually requires a separate boot or system partition.
- Windows 2000 supports RAID 0, RAID 1 and RAID 5.
- Software fault tolerance in Windows 2000 is only available on dynamic disks.

General benefits of RAID

The specific benefits of RAID really depends on the RAID solution chosen but generally, a RAID solution would provide the following two points.

1. Increased performance
Input/output (I/O) improves tremendously because several disks can be accessed simultaneously thus allowing data to be written to or read from a RAID array faster than what a single drive can offer.
2. Fault-tolerance
With redundancy built-in (except RAID 0), data stored on the array can survive the complete failure of one hard disk drive without requiring any need data restoration from backup.

Selecting a RAID level

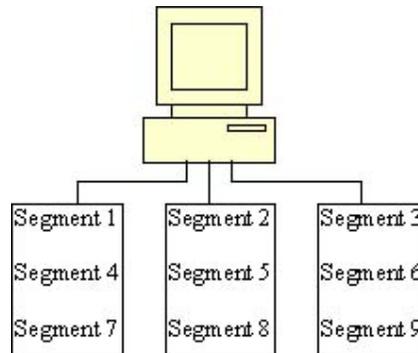
As with all things in life, there is no 'one size fits all' RAID solution. There are many levels of RAID levels available but the most commonly used ones are RAID 0, RAID 1 and RAID 5.

To achieve optimal performance, you should take into consideration the following factors:

- The number of physical disk in the RAID array
- The storage capacity of the physical disk
- Is disk performance important?
- Do you need data redundancy?
- And most importantly, does your have budget allows you to have what you want?

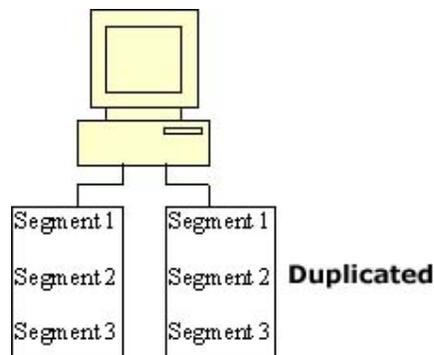
RAID 0 (Disk Striping)

- Provides disk striping across all drives in the RAID subsystem (Minimum of two to thirty two hard disk drives).
- Offer the best performance (fastest) of any RAID level BUT does not provide any data redundancy.
- Breaks up data into small blocks and then writes a block to each drive in the array. (Size of each block is determined by the stripe size parameters set during the creation of the RAID set but I recommend the use of 64KB since Windows process I/Os in 64KB).
- By breaking up a large file into smaller blocks, several drives read or write the files faster.
- Offers high bandwidth / data throughput, especially for large files as RAID 0 involves no parity calculations to complicate the write operation.
- Ideal for applications that require high bandwidth, but do not require fault tolerance.
- RAID 0 is also used to demote an 'independent' of single drive.



RAID 1 (Mirroring)

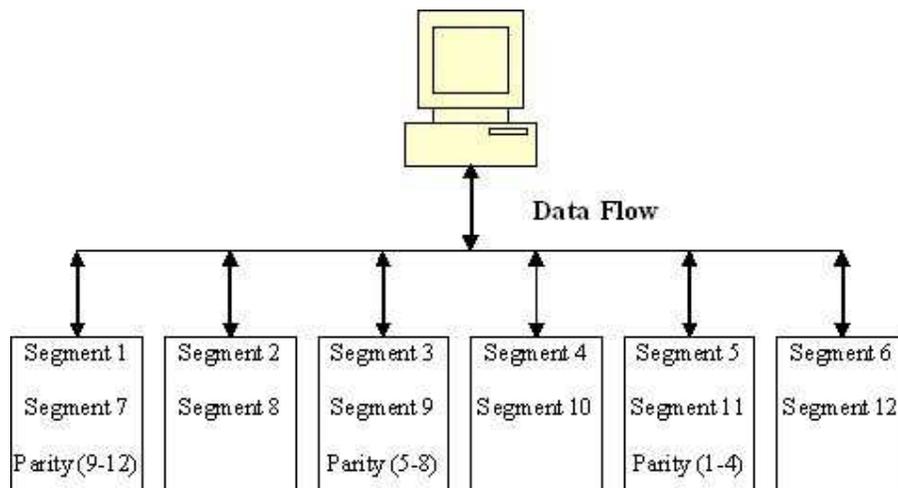
- Requires 2 hard disk drives.
- All data is duplicated from one drive to a second drive (Data is read from either drive.)
- Provides complete data redundancy BUT at cost of doubling the required data storage capacity.
- Tends to be faster on reads and slower on writes.
- Best suited for small databases or any other environment that requires fault tolerance but small capacity.
- Performance is impaired during drive rebuilds.
- Windows 2000 Server uses the “Fault Tolerance Driver” (Ftdisk.sys) to achieve software RAID 1.



Note: If the same disk controller controls both physical disks and it so happens that the disk controller fails, neither of the mirrored drives is accessible. This can be overcome by installing a second controller in the computer so that each disk in the mirrored volume has its own controller. This is known as Disk Duplexing. Disk duplexing also reduces bus traffic and improves read performance in a RAID 1 implementation.

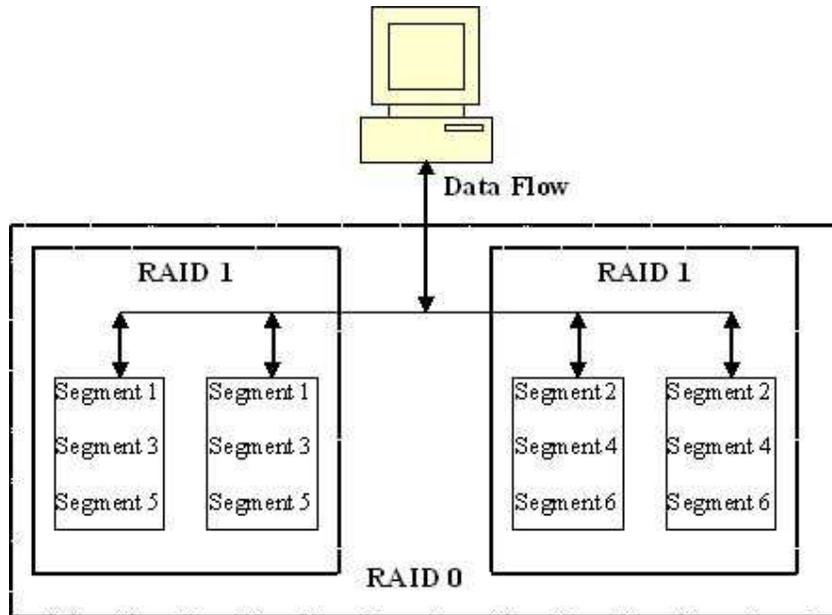
RAID 5 (Striped With Parity Volumes)

- Requires at least three, and more typically five drives are required for RAID-5 arrays.
- Includes disk striping at the block level and parity. Parity information is written to several drives.
- Addresses the bottleneck issue for random I/O operations as each drive contain both data and parity, numerous writes can take place concurrently.
- Provides high data throughput, especially for large files.
- Provides redundancy because if a drive fails, the parity drive to recreate all missing information.
- Best suited for transaction processing applications because each drive can read and write simultaneously.
- Disk drive performance will be reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
- Windows 2000 software RAID 5 requires a separate boot drive, as it cannot contain the boot or system partition.



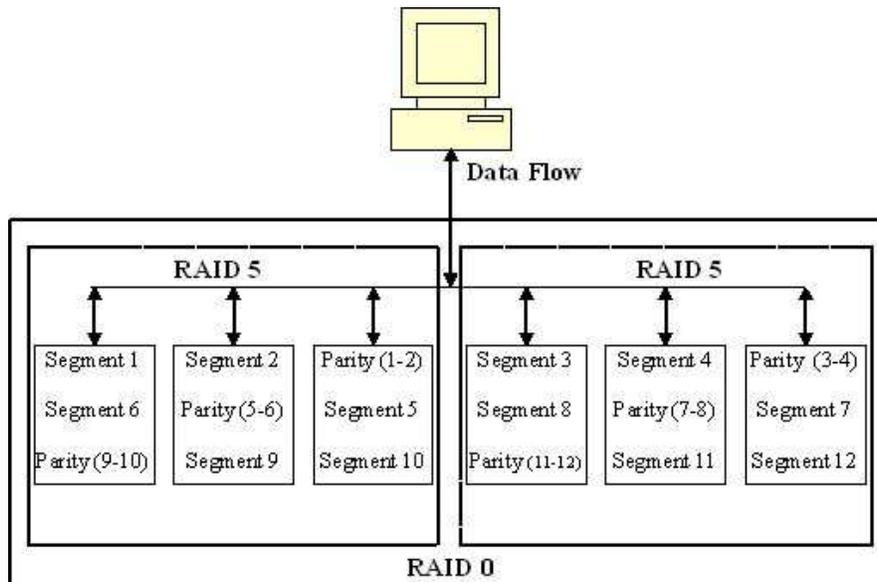
RAID 10 (RAID 0 + RAID 1)

- A combination of RAID 0 and RAID 1.
- Breaks up data into smaller blocks and then mirrors the blocks of data to each RAID 1 set.
- Able to withstand one to four hard disk drives failure if each failed disk is in a different RAID 1 array.
- RAID 10 works best for data storage that needs 100% redundancy of mirrored arrays and that also needs the enhanced I/O performance of RAID 0 (striped arrays).



RAID 50 (RAID 0 + RAID 5)

- Provides the features of both RAID 0 and RAID 5. (Parity and disk striping across multiple drives) E.g. Two RAID 5 disk arrays with data strip across both disk arrays.
- Breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then stripes the block of data to each RAID 5 array.
- Provides data redundancy and high bandwidth / data throughput for medium to large capacity.
- Requires 2 to 4 times as many parity drives as RAID 5.



Backup

While your data is relatively safer in a proper RAID implementation it is important to understand that fault-tolerant does not equate to backup. A RAID implementation with high fault-tolerance will not be useful in a disaster like a fire or a really bad day when more than one disk fails in your RAID 5 implementation. Therefore, some form of backup is still necessary as what Murphy Law states 'What ever can go wrong, will go wrong' it is wiser to be prepared than panicking when the need for data restoration is needed during an unplanned downtime.

Contact

Media and all other Queries: media@binaryhealthcare.com

About BinaryHealthcare.com

BinaryHealthcare.com is a vendor-neutral knowledge management repository pertaining to selected IT topics, Healthcare Informatics and its relevant industries (Biomedical Engineering, Radiology, Health Informatics, Telemedicine etc.) for working Professionals, students and anyone who is interested in this unique profession.

For more information, visit www.binaryhealthcare.com