IJREAT International Journal of Research in Engineering & Advanced Technology, Volume 2, Issue 5, Oct-Nov, 2014
**ISSN: 2320 – 8791 (Impact Factor: 1.479)**
**www.ijreat.org**

# SURVEY ON RAID

**Aishwarya Airen[1], Aarsh Pandit[2], Anshul Sogani[3]**

[1,2,3]A.I.T.R, Indore.

***Abstract*— RAID stands for Redundant Array of Independent Disk that is a concept which provides an efficient way for protection of data using the concepts of data striping, mirroring and parity. In this paper we have discussed about the RAID concepts and their implementations using different RAID levels. Each level of RAID is based on above concepts that are discussed and compared on the basis of their functionality, performance, cost, and reliability.**

***Index Terms*— Fault Tolerance, Hard Disk Drives, Mirroring, Parity, Performance, RAID, Solid State Drives, Striping.**

## I. INTRODUCTION

SINCE the development of computers, data storage and disk failure recovery have been one of the major challenges that have been faced. To overcome the problem RAID was introduced which is an appropriate storage and backup solution. RAID stands for "Redundant Array of Inexpensive Disks". This concept was proposed by David Patterson, Garth A. Gibson , and Randy Katz at the University of California, Berkeley in 1987. RAID was later interpreted as "Redundant Array of Independent Disks". Before this concept was introduced, people used large disks to store and backup. However, it was found that large array of disks are more prone to disk failures as compare to smaller units of disks. The probability of disk failure is equal to number of disks. Thus, this results in lower reliability. A system was required which can effectively store data efficiently and for a longer duration. Depending on requirements, some people may prioritize performance and capacity, whereas others may be more interested in security and speed. In order to meet the storage needs and maintain the backup of data, RAID is a suitable solution. The underlying basis of RAID is to take some inexpensive disks, combine them together, to make them work as one single disk. This will work as a single overall system. RAID is a data redundancy, fault tolerance, increased capacity and increased performance system. It is generally found for maintaining enormous file servers, transaction of application servers, where data accessibility and reliability is important. Hardware RAID can be embedded into a computer's motherboard or installed as an expansion card.

RAID has 6 basic levels and 2 hybrid levels. To implement a RAID system, a minimum of 2 Hard Disk Drives are necessary. Each RAID level provides different reliability, performance, efficiency and input output levels. [1][3]

## II. NEED FOR RAID

To overcome disk drive failures and data losses RAID concept was introduced. This technology has increased the protection of data and tolerance of fault by using multiple low cost (Inexpensive) disks. It increases the number of parallel disk access and to avoid the extra space for the redundant data are some of the vital factors why there is need of RAID. It comprises of various RAID levels that we will study further and their concepts starting with the first RAID level i.e. RAID 0. It uses the concept of data striping. Each RAID level is based on certain concepts. As the developments took place many more concepts were introduced such as mirroring, parity, dual-parity, distributed parity etc. the user has to decide which RAID level implementation is appropriate for him.[4][5]


## III. CONCEPTS

RAID levels are categorized on the basis of striping, mirroring and parity. They determine the data availability and performance characteristics of disk arrays. Some RAID levels use a particular technique whereas nested RAID levels use a combination of the individual techniques. Application performance and data availability requirements determine the RAID level selection. [2][3]

## MIRRORING

Mirroring or replication, as the name suggests, clearly depicts that a copy of the source exists. However, it is different from just a term copy which need not necessarily include the updating of the source when any amendments are made. Thus mirroring, as a concept can be understood as the mirror image of any source object where both source and its replica are synchronized.


Data is stored on two different HDDs (Hard Disk Drives). Thus we have two copies of data. Any changes made in the source are reflected in its mirror image. The above diagram clearly represents a mirror image of the word 'RAID'. In a similar sense, data in the disk is also mirrored. Thus a copy is created. During disk failure, it is easy to recover data by replacing it by the existing copy of the disk. This improves read performance and hence enables faster recovery from disk failure. [1][5]

## STRIPING

This concept breaks or fragments the data into blocks and spreads it across several hard disks. Stripes are blocks of a single file which is broken into smaller pieces. The set of aligned strips that spans across all the disks within the RAID set is called a stripe.

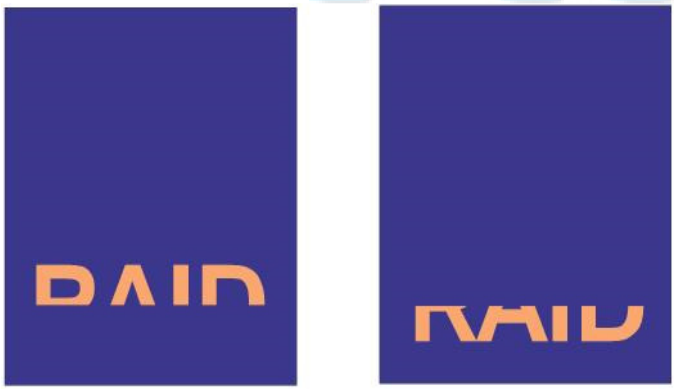

The number of blocks in a strip is strip size and is the maximum amount of data that can be written to or read from a single HDD in the set before the next HDD is accessed, assuming that the accessed data starts at the beginning of the strip. Decreasing strip size means that data is broken into smaller segments when spread across the disks. Stripe size is a multiple of strip size by the number of HDDs in the RAID set. Stripe width refers to the number of data strips in a stripe. In the above diagram, the striping technique is explained where the word 'RAID' is sliced and two stripes of the word are formed. Striped RAID does not protect data unless parity or mirroring is used. Thus it is combined with mirroring and parity to give best performance results and fault tolerance. Striping could be either byte level or block level. Depending upon the capacity of drives and need, either of the striping techniques is used. Striping improves performance by distributing data across disks. [1]

## PARITY

Parity is a protection scheme for data to resist hard disk failures. It is a logical calculation of striped data in which each striped data is XOR and stored in additional hard disk drive which stores the calculated parity of each HDD. The calculation of parity is done by the raid controller. This concept of parity takes just N+1 disk drives to store the data whereas, mirroring takes 2N disk drives to store the data. It is the concept of maintaining and recovery the data without the full set of full data which removes data redundancy which implies less disk space.

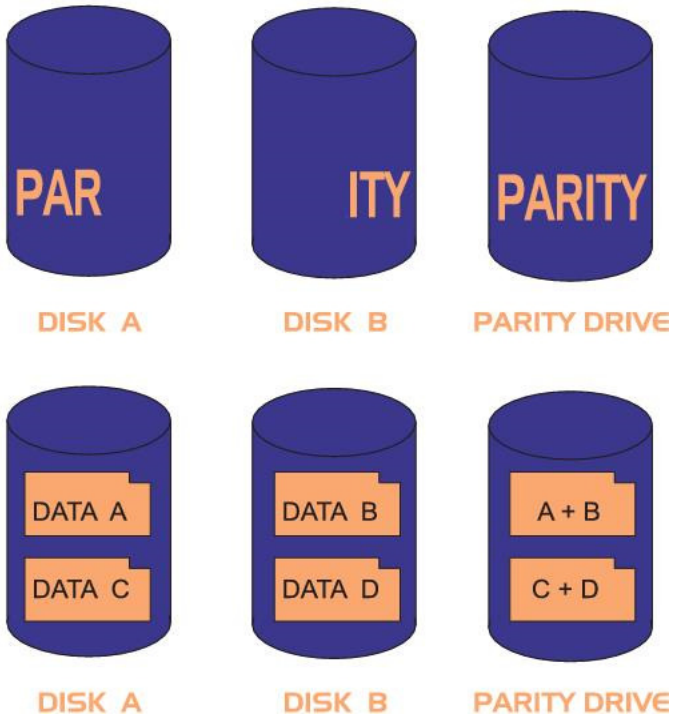

Generating parity of a data element is just like incrementing it by one and storing it in another hard drive. Suppose there are P data elements. Now we will use it to create its parity and thus its outcome will be P+1 data elements. So in any case if any one of P+1 data elements are lost, it can be recovered by knowing P data elements. The above explanation is relevant to the illustration where parity of both the disks is stored in a third disk. This parity concept has been followed in various RAID levels. RAID 3 was the first to introduce the concept of parity which is extended in RAID 4, RAID 5 and RAID 6 too as dedicated parity, distributed parity and dual distributed parity respectively. [1]

## IV. RAID LEVELS AND CONCEPT IMPLEMENTATION

### RAID 0

RAID 0 is all about speed. It totally works on the concept of striping. The data is striped across multiple disks. This RAID level contributes nothing to reliability except actually make it worse. So it involves taking two drives or more and striping the data across all of the drives. The figure illustrates how the incoming data is striped and sent across the numbered disks. This means taking all of your capacities you get to have in theory two drives double the read and write performance. If one of the drives undergoes hardware failure you will lose all of the data that is stored on both of the drives. Thus it is prone to risks and data failures and does not ensure permanent data protection. So, when it comes to reliability, this RAID level is not a good choice as it doesn't provide data protection and availability during disk failures. When the number of drives increases, the performance becomes better as more data can be read or written simultaneously. It is particularly used when high I/O throughput is required. No space is wasted provided that the hard drives used are identical. This level is more popular due to its relatively low cost and high performance gain.[1][5]

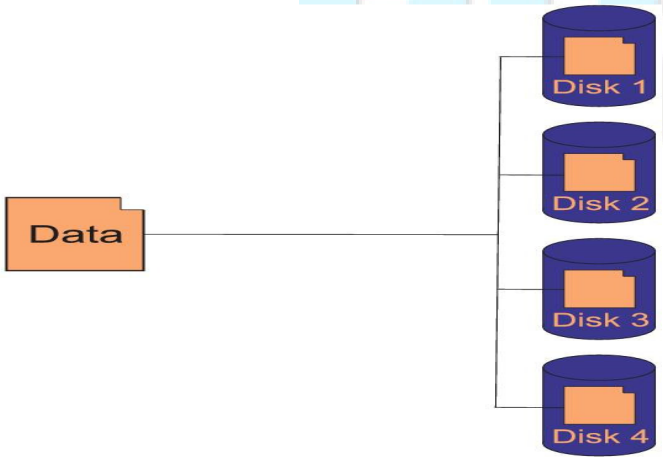


This configuration is only ideal if you are doing very frequent backups or if you are going for most extreme performance possible such as running multiple SSD's (Solid State Drive). It is also low cost as no extra storage or disk cost is required.

### RAID 1

RAID 1 is all about reliability. With the capacity of one of your drives, you get the performance of one drive but you get the redundancy of two drives. So, it consists of at least two HDD's.

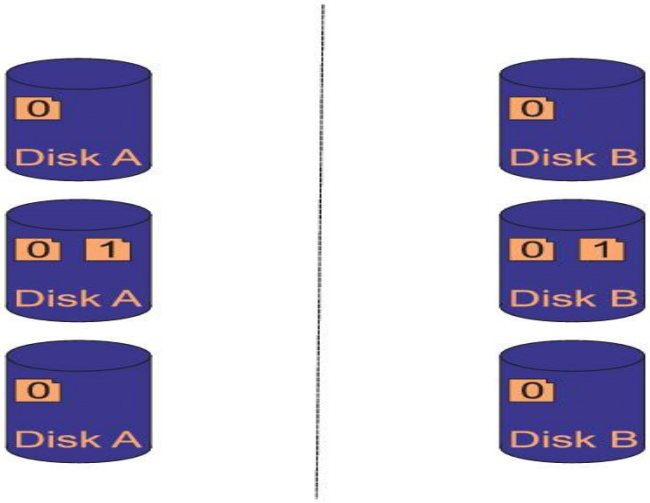


This RAID level totally works on the concept of mirroring. Here, copy of the data is stored in another disk .Every write is written to both disks. So even if one of the drives fails, the data will still be there. There is no performance overhead for running RAID 1. The more drives you add to RAID 1 you always get half the capacity than you would otherwise have. The advantage of RAID 1 is that it is extremely safe and reliable. [5]

### RAID 10/ (1+0)

This RAID level combines what is good about RAID 1 and good about RAID 0 i.e. combines mirroring and striping. For example: - Consider the following diagram. Initially mirroring of the data occurs due to implementation of the RAID 1 level concept which is mirroring. After the mirroring, the data is striped and sent across another set of disk drives. Here, implementation of RAID 0 level takes place which is the concept of striping. This concept combining RAID 0 and RAID 1 uses former for its high performance and latter for its high fault tolerance.

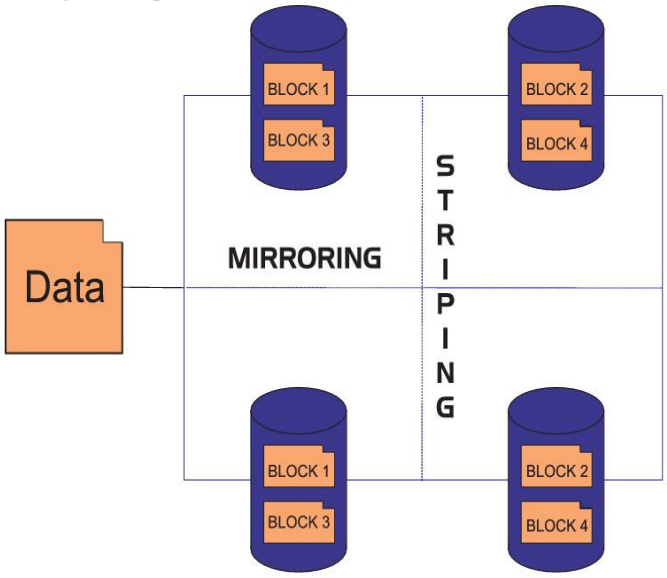


This RAID level requires minimum of four disks. We get double the performance of an individual drive. Similarly, we get double the capacity of an individual drive. But we could lose up to two drives in a RAID 10 array without losing any data. This is great where performance is needed and space is needed. [9][10]

RAID 01/(0+1)

If both speed and data security is important, then this is a better choice. It has combination of both RAID 0 and RAID 1 capabilities. This means that initially striping of data occurs followed by mirroring. This RAID level also requires minimum of four disks. This concept also combines RAID 0 and RAID 1 and uses the former for its high performance and latter for its high fault tolerance. However, RAID 01 fault tolerance is less as compared to RAID 10. The storage capacity on both RAID 01 and RAID 10 is same. Same is the case with the performance. Consider the following diagram. First split data then copy on two sets of hard drives. It will increase the speed but also keep a copy of your data for safety. In this RAID level, if any of the disks fails, the entire stripe is faulted. To recover the data, the mirroring of the data occurs where the entire stripe is copied each disk to an equivalent disk. [1][9][10]

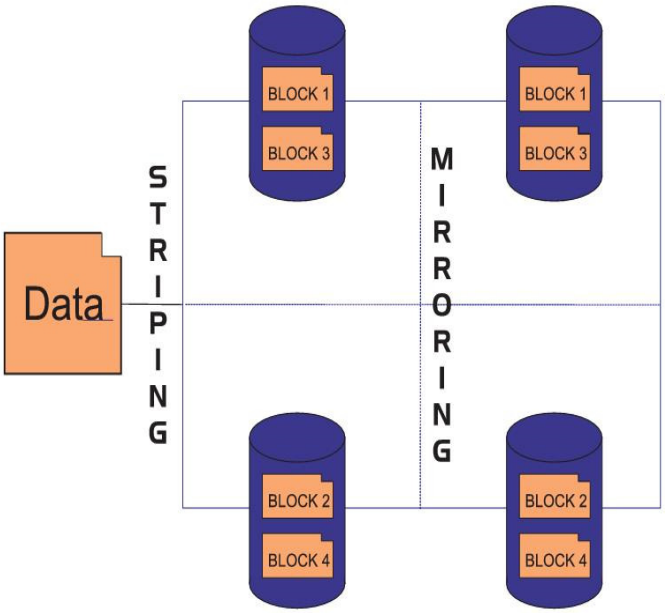


RAID 3

This RAID level uses concept of byte level striping i.e. it stripes the bits across the disks. It uses multiple data disks and dedicated parity disk. This is done so that data can be recovered if a drive undergoes failure. RAID 3 always reads and writes complete stripes of data across all the disks. However, it is inefficient for random reads and writes. This means that sequential reads and writes will have good performance. This RAID level is not commonly used. In the figure below, the parity of each block is stored in a common hard drive so that data can be recovered during disk failure. [1]

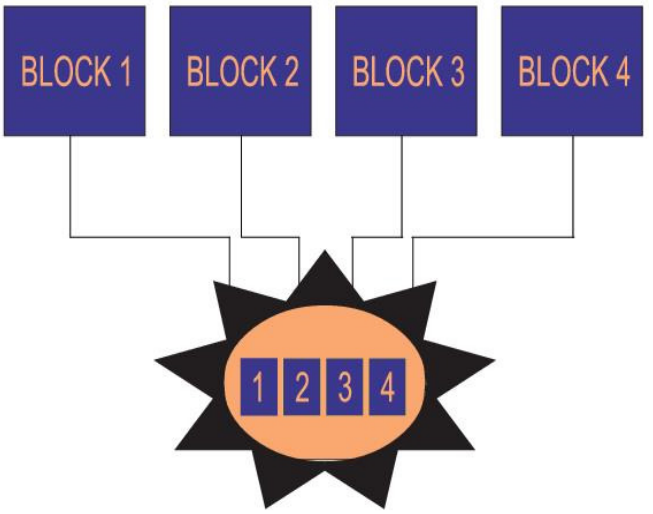


RAID 4

This RAID level uses the concept of block level striping. It uses multiple data disks and a dedicated parity disk. It has good random reads as the data blocks are striped. On the other hand, it has bad random writes, as for every write, it has to write to single parity disk. RAID 4 stripes data improving performance and uses parity for improved fault tolerance. Parity is stored in the dedicated disk so that it can be recovered at the time of failure. In the figure below, the striped blocks of data have their parity stored in a dedicated parity disk. Its cost is moderate. [1][5]

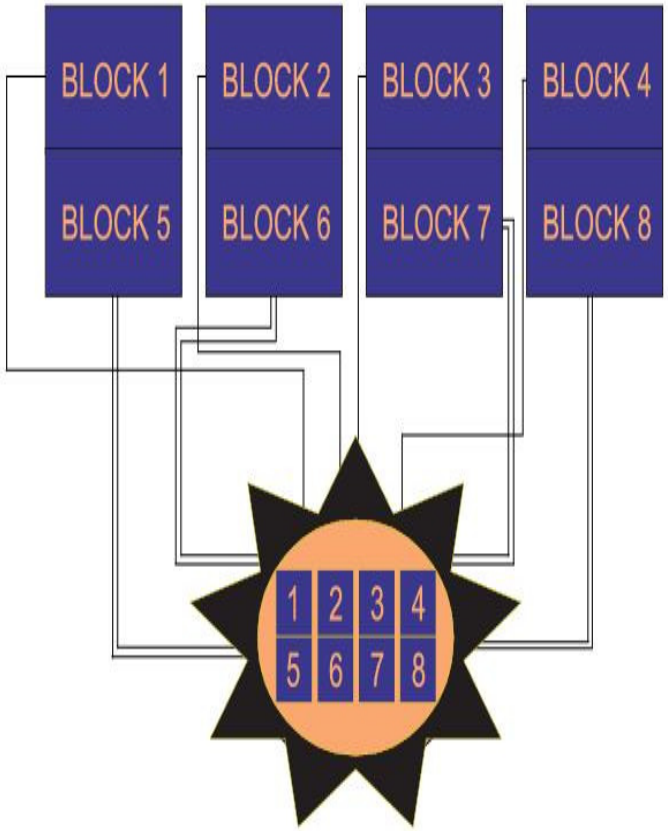


RAID 5

This RAID level is similar to RAID 4 because it also uses striping. But here the location of the parity is different. It has distributed parity and not a dedicated parity disk as in RAID 4. The utility of such a distributed parity system is that in adverse conditions when the dedicated parity disk itself fails, this parity method can help retrieve data more effectively. It is more practical for professional application and less practical for home users. Just like RAID 1, RAID 5 is for protecting your data. In an event of drive failure, it requests at least 3 drives to operate with one of the drives been reserved to rebuild the data on the array if it dies. For example, if you have 8 drives then you will have the capacity of 7 because it stores data on multiple drives .You can read from it extremely quickly making it great for archiving large amount of data. However rebuilding the array once the drive has failed and replacing it with the new one is time consuming. In the figure below, the parity of first four blocks of data are associated with block 8 and the parity of the rest four blocks are associated with block 4.[5][6]

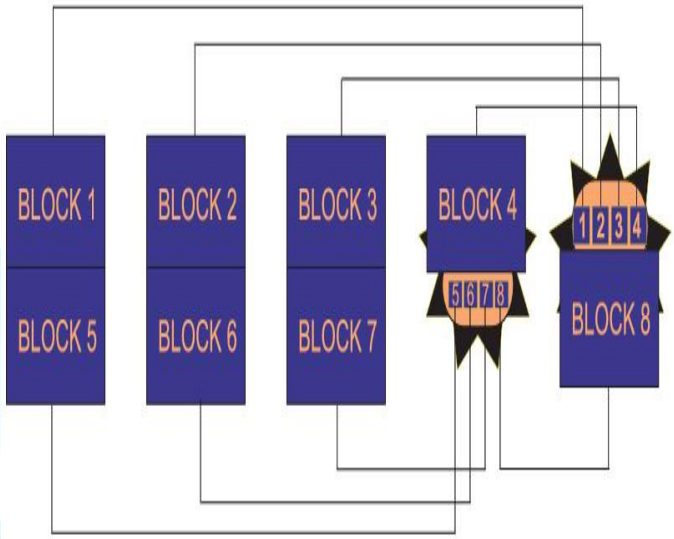


RAID 6

Just like in RAID 5 data is written across all the drives along with the parity information, similar is the case here. But this has double the parity, so the data can be retrieved. RAID 6 requires minimum of 4 drives. Usable capacity will be generally be the physical capacity less 2 drives. It is the durable version of RAID 5. It can survive up to two drive failures out of the entire array and still be completely rebuilt. However you have to have at least 4 drives and it is much slower to write than RAID. It consists of dual dedicated parity disks to enable survival in event of failure of two disks. In the figure below, the blocks have dual distributed parity. [5][6]

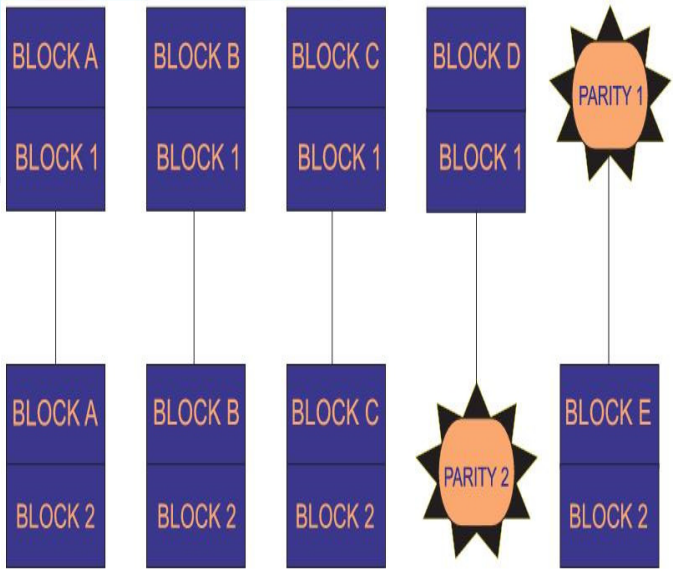

## V. COMPARING THE RAID LEVELS

Each RAID level can be compared on the basis of its characteristics like minimum number of disks required, their individual read and writes performance, applications, cost and efficiency. The purpose of this comparative study is to select the best RAID level suitable to the needs of either an individual or for an industry.

### ADVANTAGES OF RAID

RAID concepts offer advantages that increase performance and reliability of the system. Also, it makes sure that the data is retained in the system even after system crash. In general, using RAID provides data redundancy, fault tolerance, increased capacity, and increased performance. Data redundancy protects the data from hard drive failures. This benefit is good for companies or individuals that have critical or important data to protect. Fault tolerance is as important as redundancy in providing a better over-all storage system. The only RAID level that does not have any form of redundancy or fault tolerance is RAID 0. RAID also provides increased capacity by combining multiple drives. The efficiency of how the total drive storage is used depends on the RAID level. Usually, levels involving mirroring need twice as much storage to mirror the data. And lastly, the reason most people go to RAID is for the increase in performance. Depending on the RAID level used, the performance increase is different. For applications that need speed and others factors, RAID is best suitable for them. [7][8]

| Features | RAID 0 | RAID 1 | RAID 01/10 | RAID 3 | RAID 4 | RAID 5 | RAID 6 |
|---|---|---|---|---|---|---|---|
| **Min no. of disks** | 2 | 2 | 4 | 3 | 3 | 3 | 4 |
| **Data Protection** | No | 1 Disk | 1 Disk | 1 Disk | 1 Disk | 1 Disk | 2 Disk |
| **Capacity Utilization** | 100% | 50% | 50% | (n-1) *100/n | (n-1) *100/n | (n-1) *100/n | (n-2) *100/n |
| **Read Performance** | Good | Good | Very Good | Good | Very Good | Very Good | Very Good |
| **Write Performance** | Very Good | Good | Good | Poor to Fair | Poor to Fair | Fair | Good |
| **Write Penalty** | No | Average | Average | High | High | High | Very High |
| **Cost** | Low | High | High | Average | Average | Average | >RAID 5 |
| **Availability** | Low | High | Average | Average | Average | High | Very high |
| **Applications** | High end workstations, data logging, real time rendering | Operating System, Transaction Database | Fast Database, Application server | Large database ,File servers | Large files, Network storage | Data Warehousing, Web serving, Archiving | Data archive, Backup to disk, high availability solutions |

**[1][11][12][13][14]**

### DISADVANTAGES OF RAID

RAID is not a substitution for backing up. Redundancy is not the same thing as a backup. Even if you are running a RAID, you are still susceptible to viruses or accidental deletions. So make sure you do regular backups. Also, RAID does not actually make data recovery very easier. Every technology yields performance issues and so is the case with RAID. RAID doesn't always result in improved system performance. Also, not every system configuration supports RAID implementation. It needs to be made sure that the motherboard is compatible with the RAID system settings for its implementation. RAID controllers are costly. [7][8]

## CONCLUSION

Each RAID level has its own utility caused by differences in I/O processor. The reliability of redundant arrays varies from RAID levels. However, we may encounter common failures where each level may have to compromise with the performance issues. RAID is a good solution for companies or individuals demanding more transfer performance, redundancy, faster access time and storage capacity in their data storage systems. There are many levels of RAID, which range from very simple and cheap to extremely complex and expensive. The benefits of having RAID in our system are many. However, RAID is not for everyone. Performance

issues will always be there. RAID in general is so far the best method for obtaining performance, data recovery and faster access time.[4]

## REFERENCES

[1]. RAID ISM EMC Book
[2]. http://www.cs.cmu.edu/~garth/RAIDpaper/Patterson88.pdf
[3]. http://research.microsoft.com/en-us/about/our-research/default.aspx
[4]. http://www.midwestdatarecovery.com/understanding-raid-technology.html
[5]. http://www.linux-mag.com/id/7924/
[6]. http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=%2Fdiricinfo%2Ffqy0_craid1e.html
[7]. http://www.guzzzt.com/files/coding/raid.pdf
[8]. http://www.hightech-post.com/2011/06/advantages-and-disadvantages-of-raid.htmll
[9]. http://www.raid-data-recovery.net/advantages-raid.html
[10]. http://www.thegeekstuff.com/2011/10/raid10-vs-raid01/
[11]. http://aput.net/~jheiss/raid10/
[12]. http://www.raidixstorage.com/technology/raid- levels-comparison/
[13]. http://www.datarecovery.net/articles/raid-level-comparsion.html
[14]. http://www.raid-calculator.com/raid-types-reference.aspx