

UNIT IV – IMPLEMENTATION TECHNIQUES

1. RAID

(Redundant array of independent disks) originally redundant array of inexpensive disks) is a way of storing the same data in different places on multiple hard disks to protect data in the case of a drive failure.

RAID: Redundant Arrays of Independent Disks

Disk organization techniques that manage a large numbers of disks, providing a view of a single disk of high capacity and high speed by using multiple disks in parallel, and high reliability by storing data redundantly, so that data can be recovered even if a disk fails

- These levels contain the following characteristics:
- It contains a set of physical disk drives.
- In this technology, the operating system views these separate disks as a single logical disk.
- In this technology, data is distributed across the physical drives of the array.
- Redundancy disk capacity is used to store parity information.
- In case of disk failure, the parity information can be helped to recover the data.

Motivation for RAID


- Just as additional memory in form of cache, can improve the system performance, in the same way additional disks can also improve system performance.
- In RAID we can use an array of disks which operates independently since there are many disks, multiple I/O requests can be handled in parallel if the data required is on separate disks
- A single I/O operation can be handled in parallel if the data required is distributed across multiple disks.

Benefits of RAID

- Data loss can be very dangerous for an organization
- RAID technology prevents data loss due to disk failure
- RAID technology can be implemented in hardware or software
- Servers make use of RAID Technology

RAID LEVEL 0

- RAID level 0 divides data into block units and writes them across a number of disks.
- As data is placed across multiple disks. it is also called —data Striping||.
- The advantage of distributing data over disks is that if different I/O requests are pending for two different blocks of data, then there is a possibility that the requested blocks are on different disks.
- There is no parity checking of data. So if data in one drive gets corrupted then all the data would be lost.
- Thus RAID 0 does not support data recovery.
- Spanning is another term that is used with RAID level 0 because the logical disk will span all the physical drives.
- RAID 0 implementation requires minimum 2 disks.

**Advantages**


- I/O performance is greatly improved by spreading the I/O load across many channels & drives.
- Best performance is achieved when data is striped across multiple controllers with only one driver per controller

Disadvantages

- It is not fault-tolerant, failure of one drive will result in all data in an array being lost


RAID LEVEL 1: Mirroring (or shadowing)

- Also known as disk mirroring, this configuration consists of at least two drives that duplicate the storage of data. There is no striping.
- Read performance is improved since either disk can be read at the same time. Write performance is the same as for single disk storage.
- Every write is carried out on both disks. If one disk in a pair fails, data still available in the other.
- Data loss would occur only if a disk fails, and its mirror disk also fails before the system is repaired Probability of combined event is very small.



RAID LEVEL 2:

This configuration uses striping across disks, with some disks storing error checking and correcting (ECC) information. It has no advantage over RAID 3 and is no longer



- 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.

Advantages-

- On the fly' data error correction
- Extremely high data transfer rates possible
- The higher the data transfer rate required, the better the ratio of data disks to ECC disks.

Disadvantages-

- Very high ratio of ECC disks to data disks with smaller word sizes
- Entry level cost very high
- Requires very high transfer rate requirement to justify.
No commercial implementations exist

RAID LEVEL 3: Bit-Interleaved Parity

- A single parity bit is enough for error correction, not just detection, since we know which disk has failed
- When writing data, corresponding parity bits must also be computed and written to a parity bit disk
- To recover data in a damaged disk, compute XOR of bits from other disks (including parity bit disk)
- I/O operation addresses all the drives at the same time, RAID 3 cannot overlap I/O. For this reason, RAID 3 is best for single-user systems with long record applications.
- The data block is subdivided ('striped') and written on the data disks. Stripe parity is generated on Writes, recorded on the parity disk and checked on Reads.
- RAID Level 3 requires a minimum of 3 drives to implement.


Advantages

- Very high Read data transfer rate
- Very high Write data transfer rate

- Disk failure has an insignificant impact on throughput
- Low ratio of ECC (Parity) disks to data disks means high efficiency

Disadvantages


- Transaction rate equal to that of a single disk drive at best (if spindles are synchronized).
- Controller design is fairly complex.
- Very difficult and resource intensive to do as a 'software' RAID



RAID LEVEL 4: Block-Interleaved Parity

- When writing data block, corresponding block of parity bits must also be computed and written to parity disk
- To find value of a damaged block, compute XOR of bits from corresponding blocks (including parity block) from other disks.






Each entire block is written onto a data disk. Parity for same rank blocks is generated on Writes, recorded on the parity disk and checked on Reads. RAID Level 4 requires a minimum of 3 drives to implement.

Advantages

- Very high Read data transaction rate
- Low ratio of ECC (Parity) disks to data disks means high efficiency
- High aggregate Read transfer rate
- Low ratio of ECC (Parity) disks to data disks means high efficiency


Disadvantages

- Quite complex controller design
- Worst Write transaction rate and Write aggregate transfer rate
- Difficult and inefficient data rebuild in the event of disk failure
- Block Read transfer rate equal to that of a single disk



RAID LEVEL 5:

- RAID 5 uses striping as well as parity for redundancy. It is well suited for heavy read and low write operations.
- Block-Interleaved Distributed Parity; partitions data and parity among all N + 1 disks, rather than storing data in N disks and parity in 1 disk.



RAID LEVEL 6:

- This technique is similar to RAID 5, but includes a second parity scheme that is distributed across the drives in the array. The use of additional parity allows the array to continue to function even if two disks fail simultaneously. However, this extra protection comes at a cost.
- P+Q Redundancy scheme; similar to Level 5, but stores extra redundant information to guard against multiple disk failures.
- Better reliability than Level 5 at a higher cost; not used as widely.