IT8001
Information Storage and Management

Professional Elective

Topic

RAID

Varun C.M.

Assistant Professor, Department of Information Technology,
St. Xavier's Catholic College of Engineering

Expected Outcome

CO1

Understand the logical and physical components of a Storage infrastructure.

RAID

RAID (Redundant Array of Independent Disks) is an enabling technology that leverages multiple drives as part of a set that provides data protection against drive failures.

RAID Implementation Methods

Software RAID uses host-based software to provide RAID functions. It is implemented at the operating-system level and does not use a dedicated hardware controller to manage the RAID array

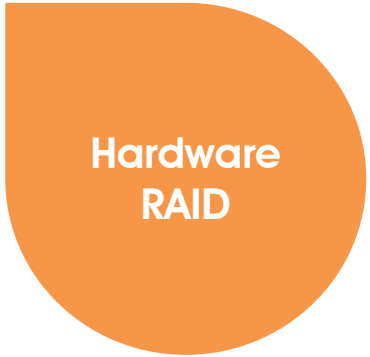

In hardware RAID implementations, a specialized hardware controller is implemented either on the host or on the array.

RAID Array Components

A RAID array is an enclosure that contains a number of disk drives and supporting hardware to implement RAID.

A subset of disks within a RAID array can be grouped to form logical associations called logical arrays, also known as a RAID set or a RAID group

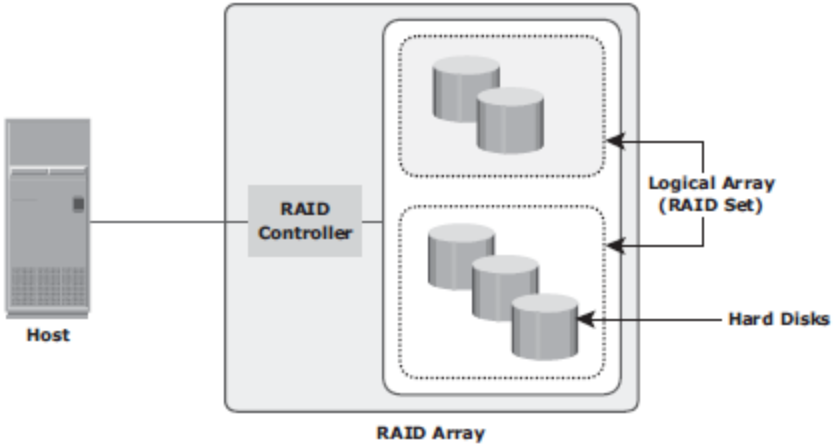

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Techniques

Striping is a technique to ***spread data across multiple drives*** to use the drives in parallel. All the read-write heads work simultaneously, allowing more data to be processed in a shorter time and increasing performance, compared to reading and writing from a single disk.

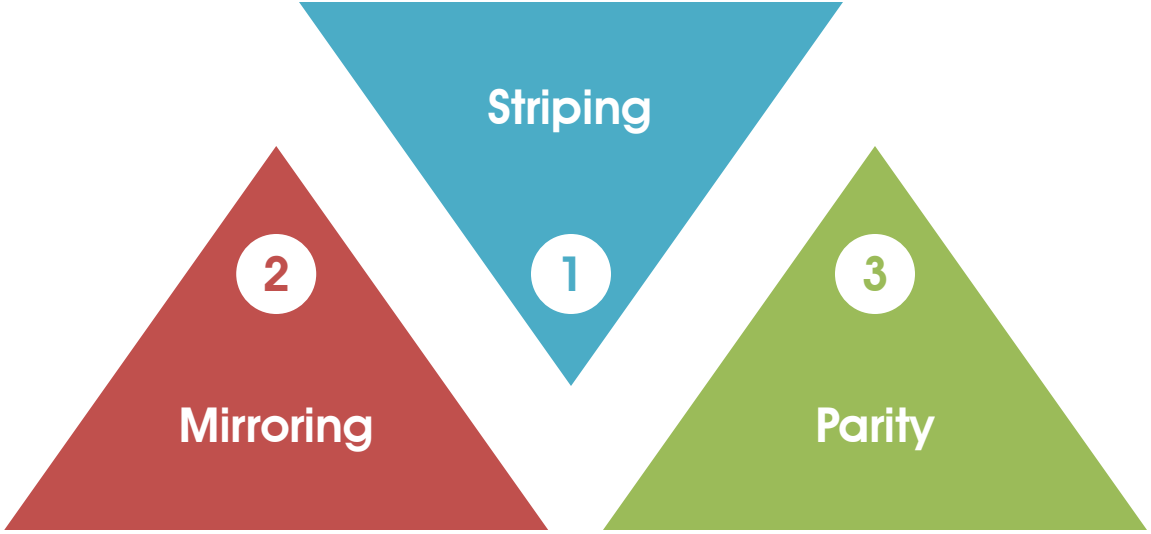

Mirroring is a technique whereby the ***same data is stored on two different disk drives***, yielding two copies of the data.

Parity is a method to protect striped data from disk drive failure without the cost of mirroring. An ***additional disk drive is added to hold parity***.

RAID Levels

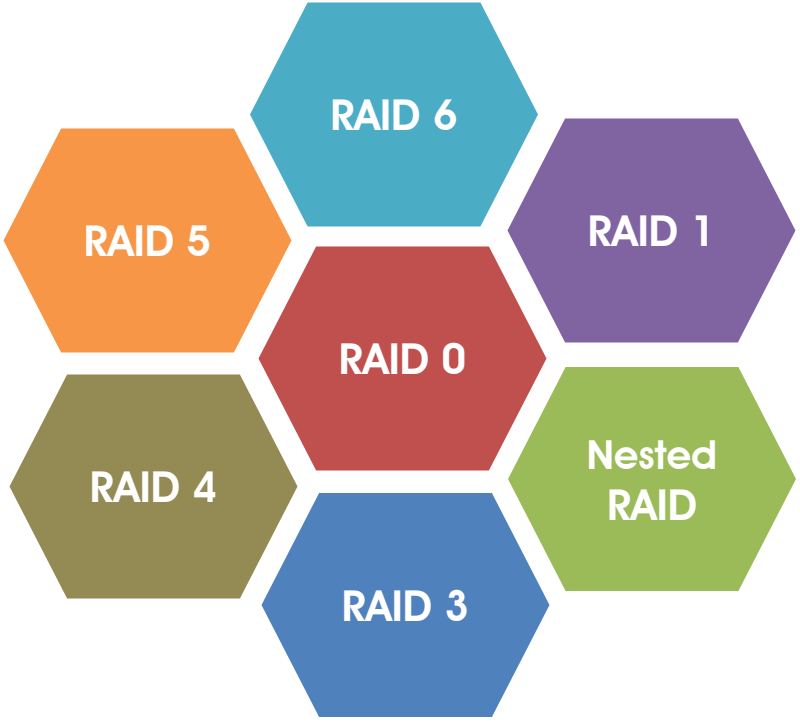

RAID Levels

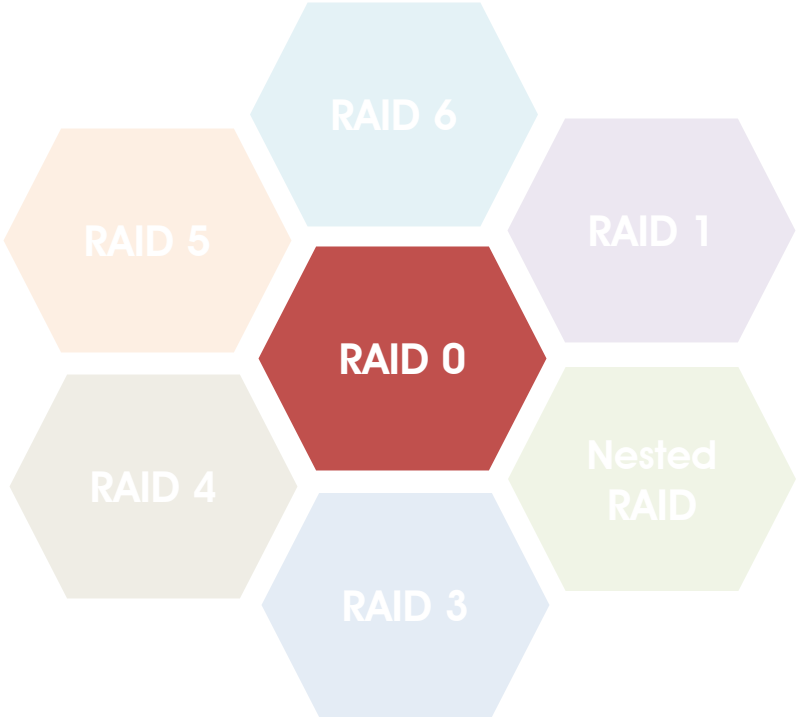

Striped set with no fault tolerance

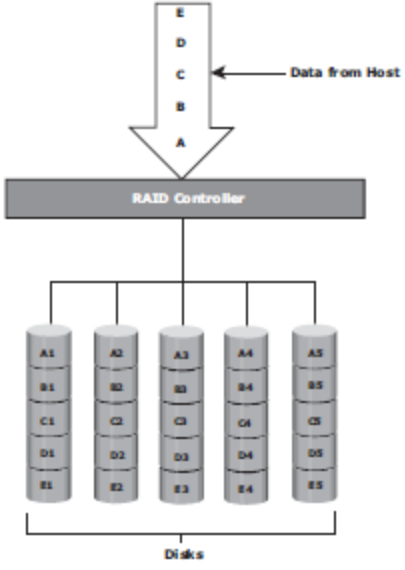

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Levels

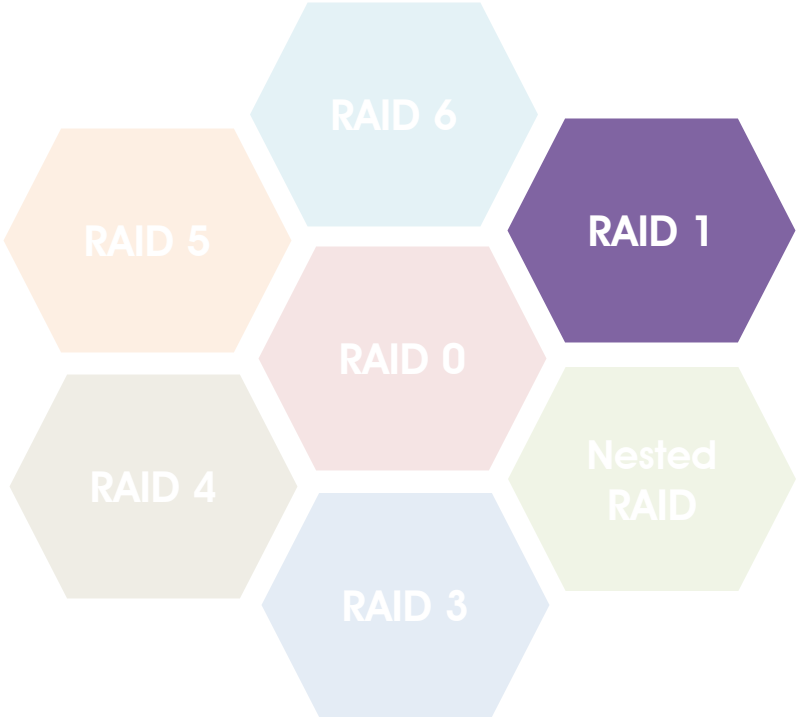

Disk mirroring

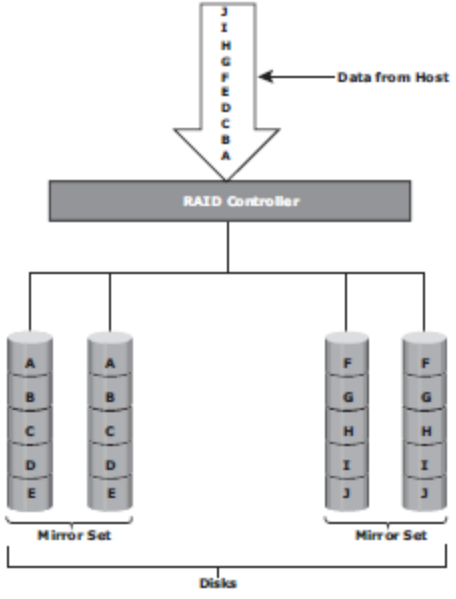

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Levels

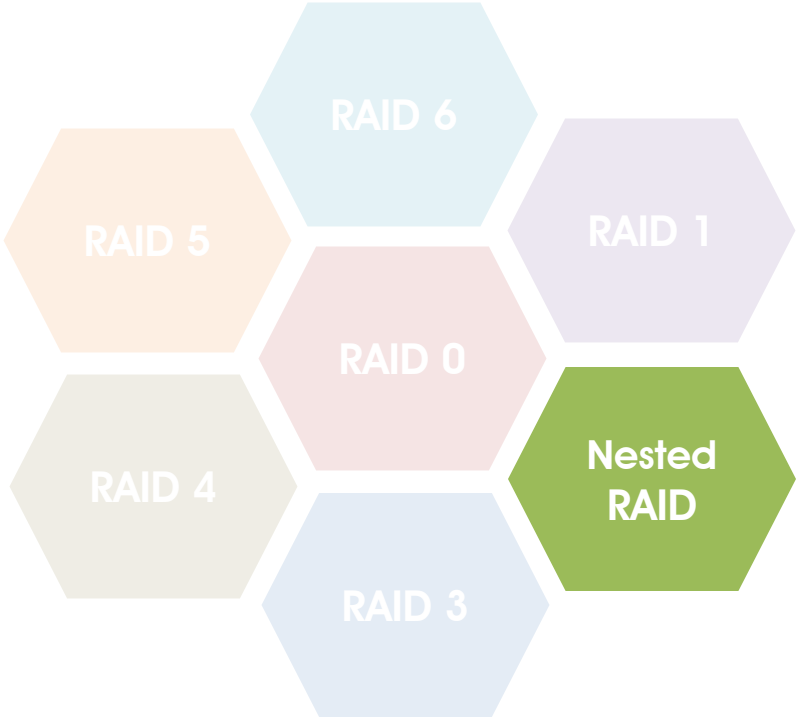

Combinations of RAID levels
RAID 1 + RAID 0

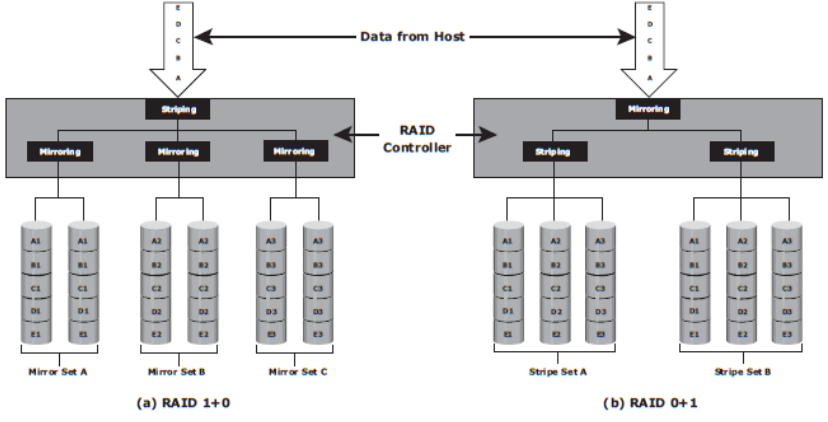

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Levels

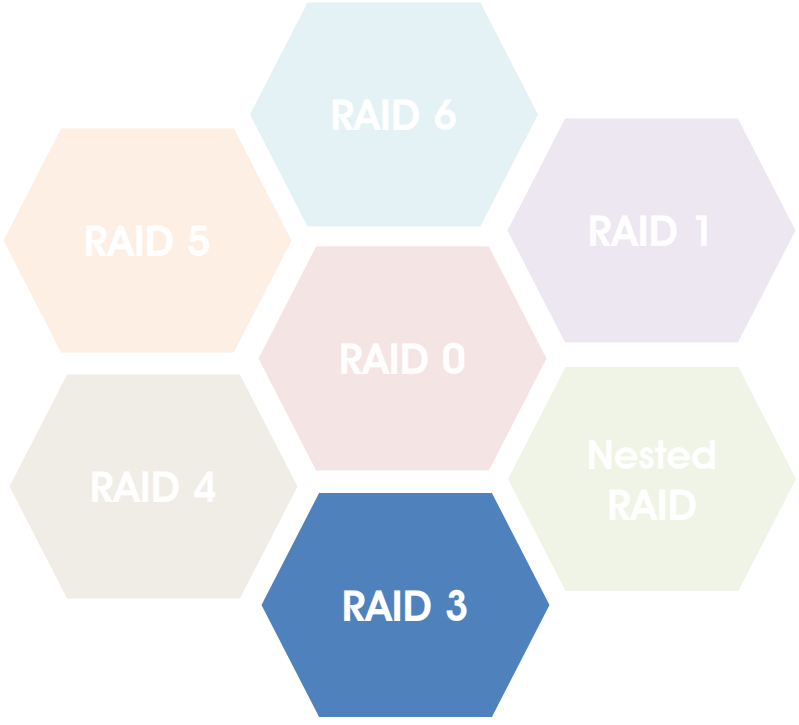

Striped set with parallel access and a dedicated parity disk

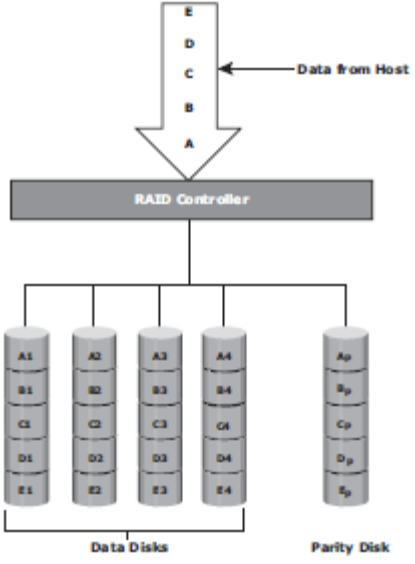

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Levels

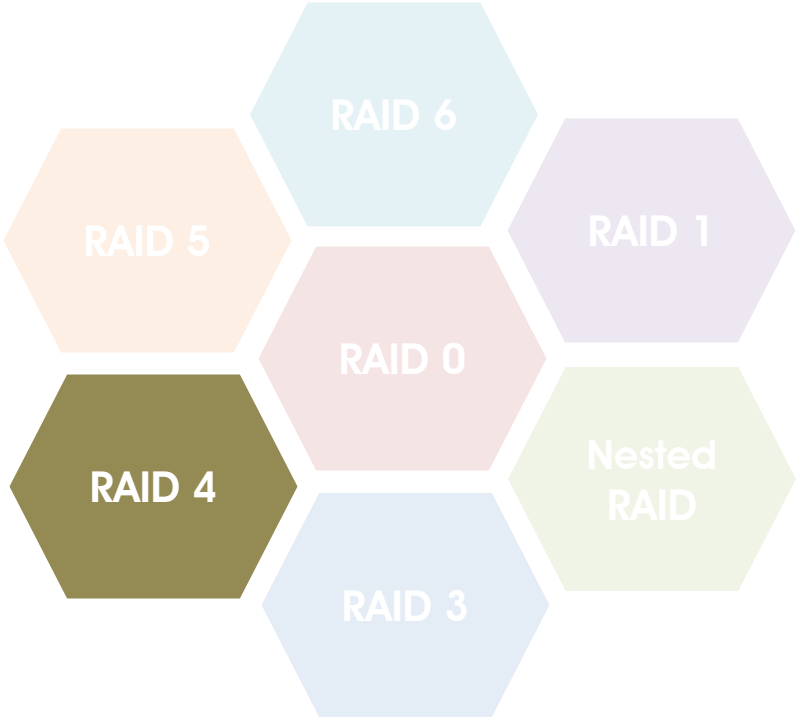

Striped set with independent disk access and a dedicated parity disk

RAID Levels

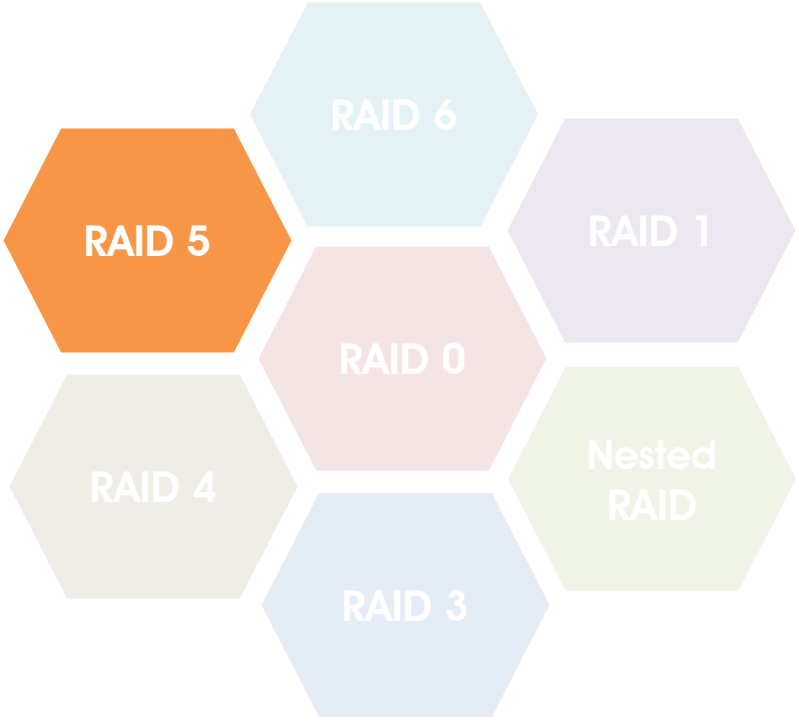

Striped set with independent disk access and distributed parity

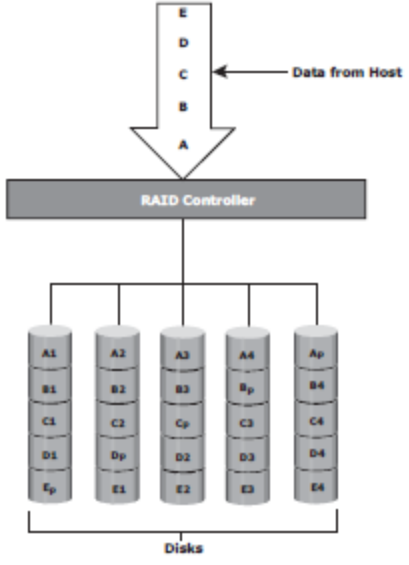

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Levels

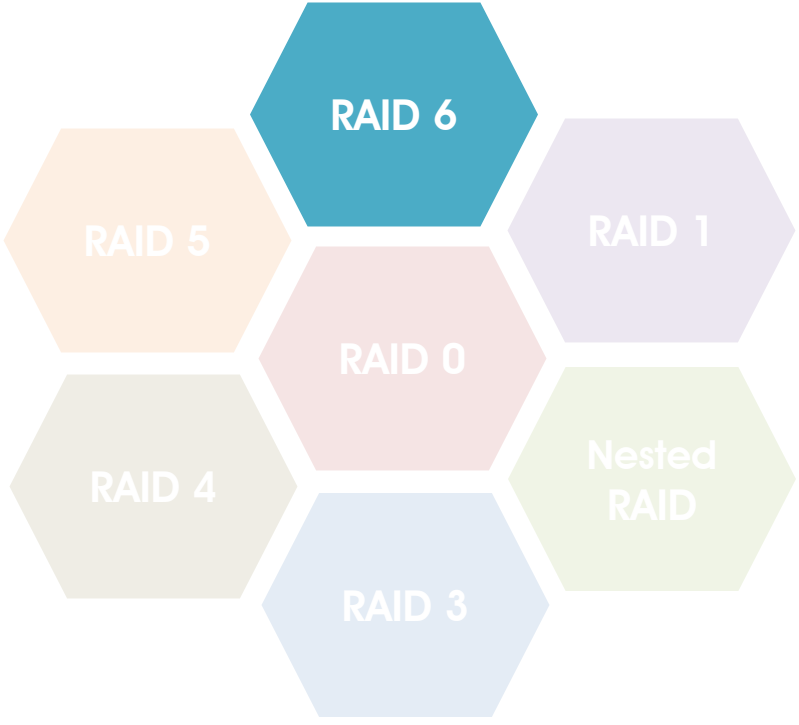

Striped set with independent disk access and dual distributed parity

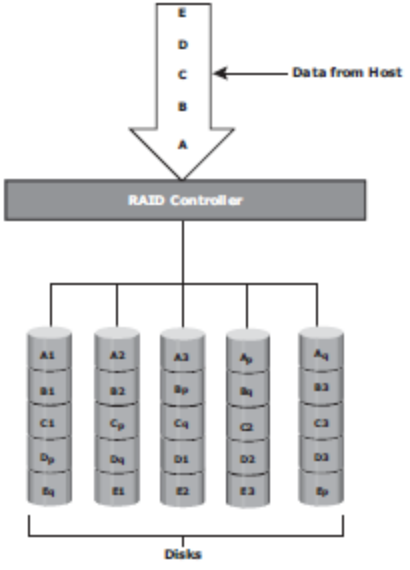

Image courtesy EMC Corporation, "Information Storage and Management"

RAID Comparison

RAID	MIN. DISKS	STORAGE EFFICIENCY %	COST	READ PERFORMANCE	WRITE PERFORMANCE	WRITE PENALTY	PROTECTION
0	2	100	Low	Good for both random and sequential reads	Good	No	No protection
1	2	50	High	Better than single disk	Slower than single disk because every write must be committed to all disks	Moderate	Mirror protection
3	3	[(n-1)/n] × 100 where n= number of disks	Moderate	Fair for random reads and good for sequential reads	Poor to fair for small random writes and fair for large, sequential writes	High	Parity protection for single disk failure
4	3	[(n-1)/n] × 100 where n= number of disks	Moderate	Good for random and sequential reads	Fair for random and sequential writes	High	Parity protection for single disk failure
5	3	[(n-1)/n] × 100 where n= number of disks	Moderate	Good for random and sequential reads	Fair for random and sequential writes	High	Parity protection for single disk failure
6	4	[(n-2)/n] × 100 where n= number of disks	Moderate but more than RAID 5.	Good for random and sequential reads	Poor to fair for random writes and fair for sequential writes	Very High	Parity protection for two disk failures
1+0 and 0+1	4	50	High	Good	Good	Moderate	Mirror protection

Image courtesy EMC Corporation, "Information Storage and Management"

Hot Spares

A hot spare refers to a spare drive in a RAID array that temporarily replaces a failed disk drive by taking the identity of the failed disk drive.

With the hot spare, one of the following methods of data recovery is performed depending on the RAID implementation

- If parity RAID is used, the data is rebuilt onto the hot spare from the parity and the data on the surviving disk drives in the RAID set.
- If mirroring is used, the data from the surviving mirror is used to copy the data onto the hot spare.

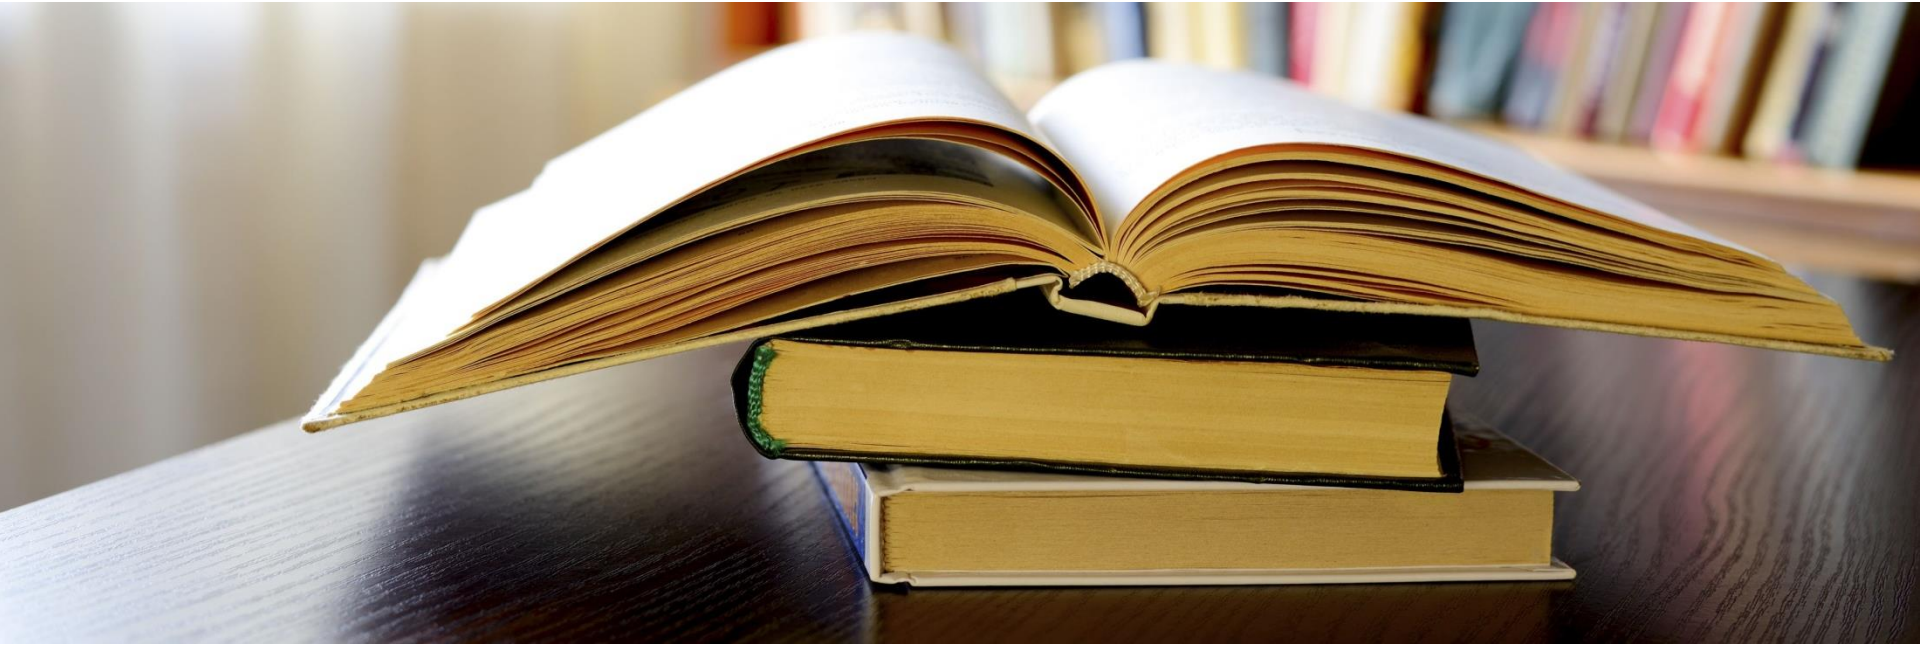

References

EMC Corporation, "Information Storage and Management" , Wiley, India.
Robert Spalding, "Storage Networks: The Complete Reference", Tata McGraw Hill, Osborne, 2003.

Varun C.M.
Assistant Professor, Department of Information Technology, St. Xavier's Catholic College of Engineering