

RAID BASICS

by David Marsh



RAID BASICS



Redundant Array of Independent Disks (RAID) is a storage technology used to improve the processing capability of storage systems. This technology is designed to provide reliability in disk array systems and to take advantage of the performance gains offered by an array of multiple disks over single-disk storage.

RAID's two primary underlying concepts are (1) that distributing data over multiple hard drives improves performance and (2) that using multiple drives properly allows for any one drive to fail without loss of data and without system downtime. In the event of a disk failure, disk access will continue normally and the failure will be transparent to the host system.

Originally designed and implemented for SCSI drives, RAID principles have been applied to SATA and SAS drives in many video systems.

RAID has been used since the 1970's but has evolved over the decades to provide performance and reliability with today's larger hard drives.

RAID Controller : is essentially, a sophisticated computer that manages data flow to and from array hard drives in the most fault-tolerant manner available.

Enclosure: The "box" which contains the controller, drives/drive trays and bays, power supplies, and fans is called an "enclosure." The enclosure includes various controls, ports, and other features used to connect the RAID to a host for example.

RAID BASICS


RAID LEVEL 0 "STRIPE"

Implements block striping where data is broken into logical blocks and striped across several drives. Although called "RAID 0," this is not a true implementation of RAID because there is no facility for redundancy. In the event of a disk failure, data is lost.

In block striping, the total disk capacity is equivalent to the sum of the capacities of all SATA/SAS drives in the array. This combination of drives appears to the system as a single logical drive.

RAID 0 provides the highest performance without redundancy. It is fast because data can be simultaneously transferred to/from multiple disks. Furthermore, read/writes to different drives can be processed concurrently.

STRIPING



No redundancy and no protection.

RAID BASICS


RAID LEVEL 1 “MIRROR”

RAID 1 implements disk mirroring where a copy of the same data is recorded onto two sets of striped drives. By keeping two copies of data on separate disks or arrays, data is protected against a disk failure. If, at any time, a disk on either side fails, the remaining good disk (copy) can provide all of the data needed, thus preventing downtime.

In disk mirroring, the total disk capacity is equivalent to half the sum of the capacities of all SAS/SATA drives in the combination. Thus, combining eight 1TB SATA drives, for example, would create a single logical drive with a total disk capacity of 4TB. This combination of drives appears to the system as a single logical drive.

RAID 1 is simple and easy to implement; however, it is more expensive as it doubles the investment required for a non-redundant disk array implementation.

In addition to the data protection RAID 1 provides, this RAID level also improves performance. In cases where multiple concurrent I/Os are occurring, these I/Os can be distributed between disk copies, thus reducing total effective data access time.



NOTE: One drawback to RAID 1 is that it does not allow running expansion. Once a RAID 1 array has been created, to expand it, the data must be backed up elsewhere before a new drive can be added. Other RAID levels permit running expansion.

RAID BASICS

RAID LEVEL 10 (1+0)

RAID 10 is technically (RAID 1 + RAID 0), a combination of RAID 1 and 0 - mirroring and striping, but without parity. RAID 10 is a stripe across a number of mirrored drives. It is implemented as a striped array whose segments are RAID 1 arrays. RAID 10 has the same fault tolerance as RAID level 1, as well as the same overhead for fault-tolerance as mirroring alone.

Advantages:


Very high I/O rates are achieved by striping RAID 1 segments.
Excellent solution for sites that would normally use RAID 1.
Great for Oracle and other databases.

Disadvantages:


Expensive to maintain.
All drives are required to move in parallel to track properly.
Limited scalability.

STRIPING

Disk 0




Disk 1




MIRROR

Disk 2



Disk 3



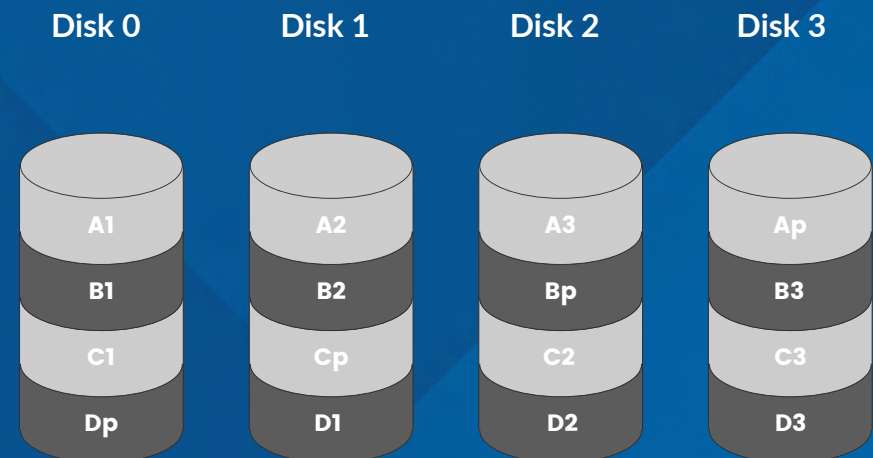
RAID BASICS

RAID LEVEL 5 "BLOCK LEVEL STRIPING"

RAID 5 implements multiple-block striping with distributed parity. This RAID level offers the same redundancy available in RAID 3; though the parity information this time is distributed across all disks in the array. Data and relative parity are never stored on the same disk. In the event a disk fails, original data can be reconstructed using the available parity information.

RAID 5 offers increased data transfer rates when data is accessed in large chunks (i.e., sequentially) and reduced data access time for many simultaneous I/O's when they do not span more than one drive.

RAID implementations include one other basic concept that needs to be introduced at this point: spare drives. RAID levels 1, 3, and 5 all allow users to include a drive as a "spare." Spare drives are installed, fully functioning, "hot-ready" hard drives which a RAID controller will use to replace a failed drive as soon as the failure is detected. The purpose of this, obviously, is to enhance the existing fault-tolerant capabilities of a RAID array.



Notice "p" parity is spread out across all of the drives in the RAID set to provide redundancy in case of hard drive failure.

RAID BASICS



RAID LEVEL 6 “BLOCK LEVEL STRIPING”


Like RAID 5, RAID 6 has multiple-block striping with distributed parity. However, double parity adds additional redundancy and allows for up to two disk failures. As drives continue to increase in size, they take much longer to rebuild which impacts performance and increases the likelihood of another failure.

(RAID 6 is a good choice when storing video with most video management software). There is no performance degradation by implementing RAID 6 over RAID 5.

While RAID 6 will cost the capacity of another drive, it's becoming essential with today's large capacity hard drives.

For more information, please contact us below.

Source: <http://www.redundantarrayofinexpensivedisks.com> and wikipedia.org



Notice “p” parity is now doubled and still spread out across all of the drives in the RAID set to provide additional redundancy in case of hard drive failure.



4660 La Jolla Village Drive, Suite 100 San Diego, CA 92122
Tel: 858-546-4471 • sales@wavereps.com • www.wavereps.com



LEARN MORE: www.wavereps.com/connect

All Material is Copyright 2018 Wave Representatives, LLC