# RAID (Redundant Array of Inexpensive Disks)

RAID stands for "Redundant Array of Inexpensive Disks" and provides a method of classifying the different ways of using **multiple** disks to increase availability and performance.
**A failed disk** (except RAID 0) **still allows users to access data on the array,** and a replacement disk or online spare can be recreated while array is in use. For RAID 3, 4, and 5, if two disks fail at the same time, all data is lost (but chance of two disk failures at once is very rare).
**SCSI** controllers are the best (and usual) implementation for RAID 0, 1, and 5. Variations of **ESDI** (Core®) or **IDE** (Compaq® IDA) are used for RAID 3 because disk synchronization is possible with ESDI and IDE.

Array = multiple physical disks treated as one logical disk.
Striping = spreading data over multiple disk drives. Data is interleaved by bytes or sectors (sector = 512 bytes) across the disks. So one file is not stored on one disk, but one file is split up and stored on multiple disks.
No RAID level allows dynamic expansion/contraction (can not change size of a defined array as storage needs change).
Hardware-based arrays *(Via hardware)* use a dedicated array processor to offload RAID functions from the CPU to usually provide better performance than software-based arrays.

## RAID 0 - Data Striping

**One file broken into sector(s) and striped across 5 disks**

*Performance:* Best - significant performance advantage over a single disk (because can do N requests simultaneously if have N disks rather than N requests to 1 disk).

*Protection:* Poor - if one disk fails, all data is lost and all disks must be reformatted (could restore array from tape however).

*Capacity:* N (N = number of disks)

*Via software:* Netware 3.x, 4.x
OASAS I for OS/2, Netware
Orthogonal Disk Array/2 (Server 195/295)

*Via hardware:* IBM RAID Controller (in 95 A and option)
IBM 3515 (High Perf Disk Subsystem)

| | | | | |
|---|---|---|---|---|
| S1 | S2 | S3 | S4 | S5 |
| S6 | S7 | S8 | S9 | S10 |
| S11 | S12 | S13 | S14 | S15 |

*Description:* data striped (spread) across each disk in array in sector(s) for improved performance.

## RAID 1 - Data Mirroring or Duplexing

*Performance:* Good - since there at least two disks, a read request can be met by either disk. Duplexing has each disk attached to a separate controller, so performance may be futher improved.

*Protection:* Good - either disk can fail and data is still accessible from other disk. With duplexing, a disk controller could fail as well and still have complete protection of data.

*Via software:* Netware 3.x, 4.x
OS/2 LAN Server - Advanced 2.0, 3.0
OASAS I for OS/2 , OASAS I for Netware
Orthogonal Disk Array/2 (Server 195/295)

*Via hardware:* IBM RAID Controller (in 95 A and option)
IBM 3515 (High Performance Disk Subsystem)

**Mirroring / Duplexing**

| | |
|---|---|
| S1 | ***S1'*** |
| S2 | ***S2'*** |
| S3 | ***S3'*** |
| S4 | ***S4'*** |

*Capacity:* N/2

*Description:* Disk mirroring duplicates data (complete file) from one disk onto a second disk using a single disk controller. **Disk duplexing** is the same as mirroring except disks are attached to a second disk controller (like two SCSI adapters).

**Hybrid RAID 1 (RAID 6)**

| | | |
|---|---|---|
| S1 | S2 | S3 |
| ***S3'*** | ***S1'*** | ***S2'*** |
| S4 | S5 | S6 |
| ***S6'*** | ***S4'*** | ***S5'*** |

Hybrid RAID 1 Array (with an odd number of disks): the first stripe is the data stripe and the second stripe is the mirror (copy) of the first data stripe but shifted one drive. *Only the IBM RAID Controller (in 95 A and option) does Hybrid RAID 1.*

## RAID 3 - Data Striping (bytes) with Parity Disk

*Performance:* Good for large transfers only - RAID 3 is generally considered better for transfer of large data blocks such as graphics or imaging files.

*Protection:* Good - if any disk fails, the data can still be accessed by using the information from the other disks and the parity disk to reconstruct it.

*Capacity:* N-1

*Via software:* None available

*Via hardware:* Third party vendors

**Data disks** | **Parity disk**

| | | | | |
|---|---|---|---|---|
| B1 | B2 | B3 | B4 | ***P1-4*** |
| B5 | B6 | B7 | B8 | ***P5-8*** |
| B9 | B10 | B11 | B12 | ***P9-12*** |
| B13 | B14 | B15 | B16 | ***P13-16*** |
| B17 | B18 | B19 | B20 | ***P17-20*** |

(B1 through B20 represent Bytes of data from one file. P1 through P20 represent the Parity of Bytes B1 to B20)

*Description:* RAID 3 stripes data, one byte at a time, across all the data drives. Parity information, used to reconstruct missing data, is stored on a dedicated drive. RAID 3 requires at least two data disks, but works best with four disks (and one parity disk).

**RAID 4** - Similar to RAID 3 with interleaved sectors (not bytes) and unsynchronized rotation (not synchronized rotation).

## RAID 5 - Data/Parity Striping (sectors)

*Performance:* Good for networks - RAID 5 is preferred for smaller block transfers the size of typical network files.

*Protection:* Good - if any disk fails, the data can still be accessed by using the info from the other disks along with the stripped parity info.

*Via software:*
OASAS I for OS/2 , OASAS I for Netware
Orthogonal Disk Array/2 (Server 195/295)

*Capacity:* N-1

*Via hardware:*
IBM RAID Controller (in 95 A and option)
IBM 3514 (High Avail Ext Disk Array)

*Description:* Raid 5 stripes data, sector(s) at a time, across all disks. Parity is interleaved with data information rather than stored on a dedicated drive. RAID 5 works with a minimum of three disks.

**Physical disks for data and parity** (S1 through S20 represent Sectors of data from one file. P1 through P20 represent the parity of sectors S1 to S20)

| | | | | |
|---|---|---|---|---|
| ***P1-4*** | S1 | S2 | S3 | S4 |
| S5 | ***P5-8*** | S6 | S7 | S8 |
| S9 | S10 | ***P9-12*** | S11 | S12 |
| S13 | S14 | S15 | ***P13-16*** | S16 |
| S17 | S18 | S19 | S20 | ***P17-20*** |

## Orthogonal RAID 5 - Data/Parity Striping (sectors) with redundant disk controllers

*Performance:* Good - same as RAID 5 (read above). Improved performance due to load being spread across disk adapters if multiple disks per controller.

*Protection:* Best - same protection as RAID 5 (above) with additional protection from any single disk controller failure.

*Capacity:* N-1

*Via software:*
OASAS I for OS/2 , OASAS I for Netware
Orthogonal Disk Array/2 (Server 195/295)

*Via hardware:* None available

*Description:* Same as RAID 5 (read above) with the grouping of disks "orthogonally" (right angle) to disk controllers, SCSI buses, and cables. So each disk has its own disk controller, bus, and cable.

| Disk Cont. | Disk Cont. | Disk Cont. | Disk Cont. |
|---|---|---|---|
| ***P1-3*** | S1 | S2 | S3 |
| S4 | ***P4-6*** | S5 | S6 |
| S7 | S8 | ***P7-9*** | S9 |
| S10 | S11 | S12 | ***P10-12*** |

In any RAID 3, 4, or 5 array, the equivalent of one disk is consumed by parity info.
Other third party vendors have software and hardware implementations of most RAID levels.
RAID 5 degrades throughput of a server about 35% compared to RAID 0.



(RAID) Compiled by Roger Dodson, IBM. July 1994