HEWLETT® PACKARD

# technology brief

## RAID Levels

**March 1997**

### Introduction

*RAID* is an acronym for *Redundant Array of Independent Disks* (originally *Redundant Array of Inexpensive Disks*) coined in a 1987 University of California at Berkeley paper by Patterson, Gibson, and Katz. RAID disk array techniques allow sets of drives to be grouped for availability purposes, providing the redundancy necessary for data protection and sometimes providing a performance advantage as well.

This paper describes seven RAID levels—0, 1, 3, 5, 10, 30, and 50. The first four (0, 1, 3, and 5) are established industry standards. Levels 10, 30, and 50 reflect the functionality provided with the HP NetRAID array controller. Knowing the characteristics of each level will help you decide which level best suits your network requirements. A section at the end of this paper provides guidelines to help you choose the level best suited to your needs.

It is important to note that RAID levels can be implemented with software or hardware. Many, but not all, network operating systems support at least some RAID levels up to level 5. Levels 10, 30, and 50 are only possible with a disk array controller (DAC). Software-based RAID uses host CPU cycles and system memory, thereby adding overhead that may impact system performance. DACs move the burden of RAID computations and manipulation from software to specialized hardware and often perform better than software RAID.

### Characteristics of RAID Levels

**RAID-0** RAID-0 (see Figure 1) uses a technique called *striping* to distribute data across disks, with stripes divided into contiguous pieces called *blocks*. Striping allows data to be accessed on multiple disks simultaneously, balancing the I/O load across the disks and resulting in maximum data capacity and fast access time. RAID-0 is the only RAID level that is not redundant. This lack of redundancy adds a low-cost advantage to the speed of RAID-0, but it also means that if one hard disk in an array fails, all data in the array will be lost. Recovering from a disk failure with RAID-0 requires replacing the failed drive and restoring data to all drives from a backup.

For networks that can afford the operating time lost to restore operations when a disk fails, RAID-0 offers a high-performance alternative. It is available with both software and hardware RAID.

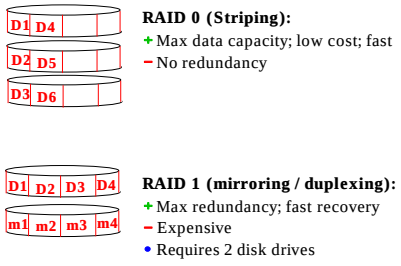


**Figure 1. RAID-0 and RAID-1**

**RAID-1** RAID-1 (see Figure 1) is also known as mirroring because the data on one disk is completely duplicated on another disk. If one fails, the other remains available so data loss and system interruptions due to disk drive failure are virtually eliminated. Duplexing is another data protection technique possible with RAID-1. While mirroring duplicates each disk, duplexing duplicates the disk controller as well, protecting your data in case of either a disk or controller failure. The disadvantage of mirroring and duplexing is the high cost of duplicating each drive or controller, which can become a serious cost issue in large servers. RAID-1 is available with both software and hardware RAID.

**RAID-3** With RAID-3 (see Figure 2), also known as striping with dedicated parity, each stripe has the equivalent of one block used for data redundancy information, or *parity*. The parity is encoded information that can be used to reconstruct the data on that stripe if one of the disks fails.

The HP NetRAID controller version of RAID-3 is striped data with a dedicated parity drive. It differs physically from the original definition of RAID-3, which includes synchronized spindles and a separate channel for each drive. However, the resulting performance of the HP NetRAID controller version of RAID-3 is similar to the original. Firmware in the controller optimizes RAID-3 data flow for long, serial data transfers. Typical environments for such transfers are video or imaging applications.

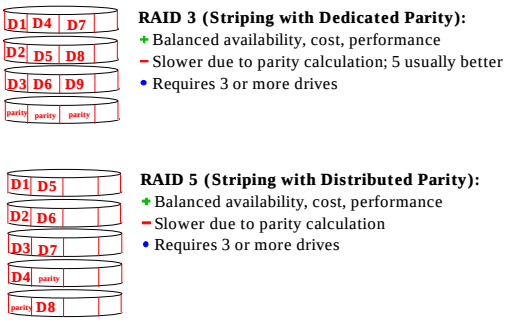


**Figure 2. RAID-3 and RAID-5**

**RAID-5** With RAID-5 (see Figure 2), also known as striping with distributed parity, each stripe has the equivalent of one block used for parity. Unlike RAID-3, RAID 5 distributes this parity equally among all disk drives. RAID-5 is the most commonly used data protection scheme because it provides good overall performance with a minimum loss in capacity. It is suitable for I/O intensive, high read/write ratio applications such as transaction processing.

To provide RAID-5 redundancy, a disk array needs a minimum of three disks (excluding a hot spare). RAID-5 is available through DAC hardware and with some network operating system software.

**RAID-10** With RAID-10 (see Figure 3), also known as striping of mirrored arrays, data is striped across disks, as in RAID-0, and each disk has a mirror disk, as in RAID-1. RAID-10 provides 100% data redundancy and supports larger volume sizes, but at a relatively high cost. RAID-10 provides the best performance for most applications where redundancy is required and cost is not a factor. Higher reliability can be achieved with RAID-10 because two physical disk drive failures (one in each array) can occur and data will still be protected. RAID-10 requires four drives and is available only with the HP NetRAID DAC.

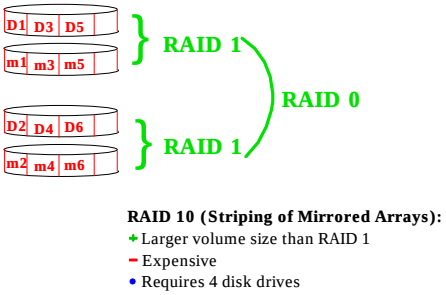


**Figure 3. RAID-10**

**RAID-30** With RAID-30 (see Figure 4), also known as striping of dedicated parity arrays, data is striped across disks, as in RAID-0, and uses dedicated parity, as in RAID-3. RAID-30 provides high fault tolerance and supports larger volume sizes. As with RAID-10, RAID-30 provides high reliability because even if two physical disk drives (one in each array) fail, data is still available.

RAID-30 requires a minimum of six drives and is available only with the HP NetRAID DAC. It is best for non-interactive applications, such as video streams, graphics, and imaging that process large files sequentially and require high availability and high speed.

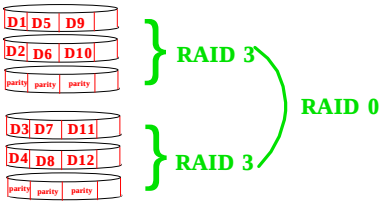


**RAID 30 (Striping of Dedicated Parity Arrays):**
+ Larger volume than RAID 3; performance in read environments
− (similar to RAID 3)
• Requires 6, 8, 10, 12, 14, or 16 drives

**Figure 4. RAID-30**

**RAID-50** With RAID-50 (see Figure 5), also known as striping of distributed parity arrays, data is striped across disks, as in RAID-0, and uses distributed parity, as in RAID-5. RAID-50 provides data reliability, good overall performance, and supports larger volume sizes. As with RAID-10 and RAID-30, even if two physical disk drives (one in each array) fail with RAID-50, no data will be lost.

RAID-50 requires a minimum of six drives and is available only with the HP NetRAID DAC. It is best for applications that require highly reliable storage, high request (read) rates, and high data transfer performance, such as transaction processing and office applications with many users accessing small files.

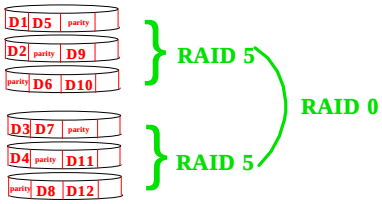


**RAID 50 (Striping of Distributed Parity Arrays):**
+ Larger volume than RAID 5; performance in read environments
− (similar to RAID 5)
• Requires 6, 8, 10, 12, 14, or 16 drives

**Figure 5. RAID-50**

Table 1 summarizes the characteristics of each RAID level.

**Table 1. RAID Level Characteristics**

| **RAID Level*** | RAID-0 | RAID-1 | RAID-3 | RAID-5 | RAID-10 | RAID-30 | RAID-50 |
|---|---|---|---|---|---|---|---|
| **Also Known As** | Striping | Mirroring | Striping with dedicated parity | Striping with distributed parity | Striping of mirrored arrays | Striping of dedicated parity arrays | Striping of distributed parity arrays |
| **Fault Tolerance** | No | Yes | Yes | Yes | Yes | Yes | Yes |
| **Redundancy Type** | None | Duplicate | Parity | Parity | Duplicate | Parity | Parity |
| **Hot Spare Option** | No | Yes | Yes | Yes | Yes | Yes | Yes |
| **Disks Required** | One or more | Two only | Three or more | Three or more | Four only | 6, 8, 10, 12, 14, or 16 | 6, 8, 10, 12, 14, or 16 |
| **Usable Capacity** | Total capacity of the disk(s) | 50% of the capacity of the disk(s) | n-1/n of the total capacity of the disk(s), where n=the number of disks | n-1/n of the total capacity of the disk(s), where n=the number of disks | 50% of the capacity of the disk(s) | n-2/n of the total capacity of the disk(s), where n=the number of disks | n-2/n of the total , capacity of the disk(s), where n=the number of disks |

*With the HP NetRAID controller, the RAID level is specified for each logical drive. Up to eight logical drives may be defined for each controller.

## Choosing a RAID Level

Three main factors will influence your choice of RAID levels: availability (data redundancy), performance, and cost. If availability is not a requirement, choose RAID-0 for maximum performance. If both availability and performance are important and cost is not a major concern, choose RAID-1 or RAID-10, depending on the number of disk drives. If cost, availability, and performance are equally important, choose RAID-3, RAID-30, RAID-5, or RAID-50, depending on the typical data transfers and the number of disk drives.

The flowchart in Figure 6 provides some guidelines for selecting a RAID level. Use it to help you determine the RAID level you need, keeping in mind that the unique characteristics of your applications might make a different RAID level more appropriate for your environment.

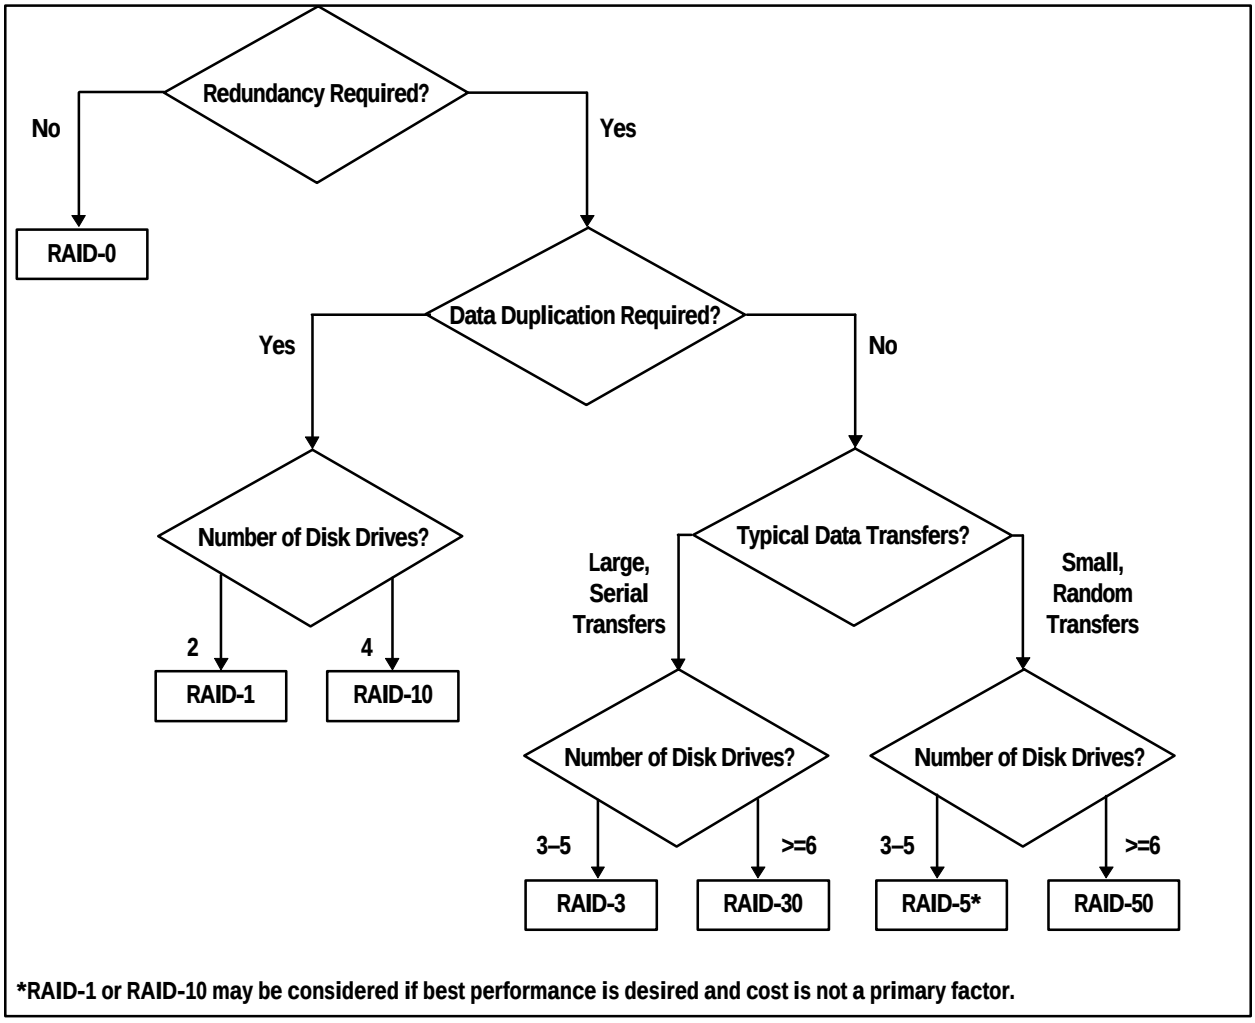


**Figure 6. Flowchart for Choosing a RAID Level**

Technical information in this document is subject to change without notice.