
Understanding RAID3

RAID3 is not supported on your XP disk array, but it is a preface to understanding RAID5, which is supported.

RAID3 allows data to be split and distributed on two or more disk drives using striping. Because two or more disk drives are recording the data simultaneously, the transfer speed is faster. The multiple disk drives that are used to store the data are called a parity group. See figure 14 (page 98).

RAID3 stores parity information on a separate drive (called a parity drive). This improves the safety of your data because the parity information can be used to recover the data, even if one of the drives in a parity group has become inoperative or has read errors.

Because RAID3 does not have parallel input/output capability, it must drive two or more drives at the same time. This causes poor performance in applications that process small data files often (transaction processing). RAID3 has good performance when you are using applications that process large data files in a single run (scientific computations).

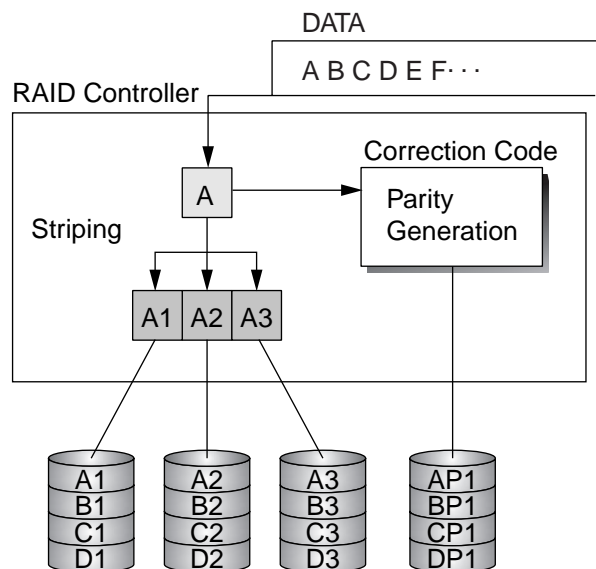


Figure 14. RAID3 Disk Array

Using RAID5

RAID5, like RAID3, processes data between two or more disk drives simultaneously. RAID5 also saves parity information to a separate disk drive.

RAID5, unlike RAID3, saves information as blocks of data, which allows the RAID controller to access each disk for only one stripe of data. This also allows RAID5 to perform input/output operations on other disks in parallel. By saving the data as small blocks, this approach permits increased I/O performance. This makes RAID5 ideal for transaction processing.

In large-scale (sequential) input/output operations, RAID5 permits parallel processing of the parity blocks, which increases the data transfer rate.

In small-scale input/output operations though, RAID5 must perform extra read operations from the data and parity disks. This slows the transfer rate and is called write penalty. To limit this problem, RAID5 distributes the parity data on several disks in the parity group.

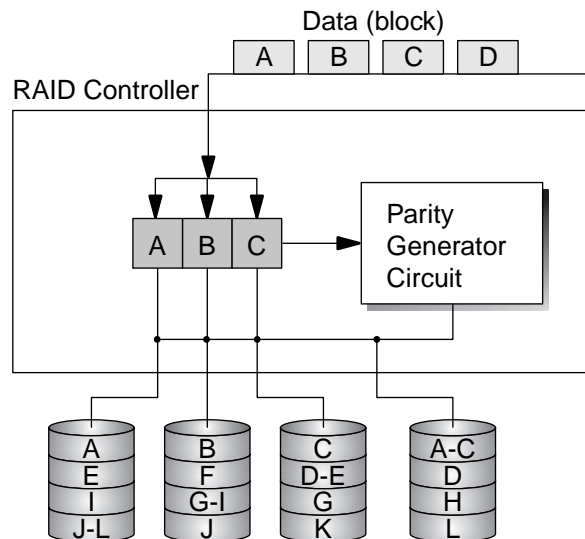


Figure 15. RAID5 Disk Array

RAID5 requires four disks for an array group. RAID5 on the XP512 allows the data to be split and distributed onto three disk drives using striping. Parity data for the group is created and stored on the parity disk as in figure 16 (page 101). Data can easily be recovered if a device in the parity group becomes inoperative or causes a read error.

In RAID5, the striping size is set to that of blocks that are to be transferred. This allows the RAID controller to access each disk for a single stripe equivalence of data and allows the RAID controller to perform I/O operations on other disks in parallel; therefore, increasing I/O performance substantially.

In small-scale or random I/O applications, the data transfer rate remains the same as conventional RAID systems. In large or sequential I/O applications, RAID5 permits the blocks in the same parity group to be processed in parallel, resulting in an increase in the data transfer rate. For small writes to single blocks, RAID5 requires extra reads from the data and parity disks before a small block can be written to the disk. Since the parity data is distributed on all disks in the group, it still allows parallel I/O processing of multiple blocks.

If the parity disk were fixed at a single disk device, the parity disk would be busied during a single write that is executed to update the parity data. This would make it impossible to perform parallel I/O processing because the parity disk would always be busy with the first block of data. This is a problem with RAID3 that RAID5 has alleviated.