

A Little Bit About RAID Arrays

By Terry Irwin

Ok, so you're looking to purchase a new server for your business and your IT department (or your IT professional) just got really excited and started speaking a language you don't understand. They were particularly exuberant about the RAID5 configuration. All the while you're thinking to yourself, "What on earth is RAID5?" I'm glad you asked.

RAID stands for Redundant Array of Inexpensive Disks. The point of RAID is to either improve performance, or have some form of fail-safe storage of data, or both. Here are three types of RAID that you are most likely to encounter as a small business owner or manager.

RAID1 (otherwise known as mirroring) is basically this: The RAID controller (a device in the computer that you connect the hard drives to, and the operating system accesses the hard drives through) uses two identical hard drives but only lets the operating system know about one of them. Behind the scenes, every operation that it does to one drive, it also does to the other drive, leaving you with 2 identical drives. That way, if one should fail, all your data is still on the other drive. If you use two 80 GB hard drives, the Operating System will only see one 80 GB hard drive. This gives you a good fail-safe system for your data but can negatively impact the performance of the system, and it costs you twice as much as just using a single hard drive.

RAID0 (otherwise known as striping) is very similar, but rather than protecting your data, it gives you a performance boost. Just like a mirrored set, you need at least 2 drives connected to the RAID controller, and it also reports to the operating system that there's only one drive there. But instead of hiding the second drive, it combines the drives to make them appear to be a single drive that's twice as big. Every time the operating system wants to write a file, the RAID controller writes part of the file to each drive because the operating system can send data to the RAID controller faster than either drive can carry out the write operation. This way, the RAID controller can write twice as much data at a time. The big draw back to this type of RAID is that you are doubling your chances of losing your data. If either hard drive fails, the whole thing is unusable. You can't get your data off of the other drive because only parts of it were stored there in the first place. If you use two 80 GB hard drives, the Operating System will only see one 160 GB hard drive.

Now what if you want the performance of a striped set, with the high data protection of a striped set? That's where advanced forms of RAID kick in, like RAID5. RAID5 is striping with parity. To understand why this protects your data, you have to first understand parity. Parity is a way of tracking not the data itself, but information ABOUT the data. I'll write about parity in another document at some point in the future.

RAID5 is very similar to RAID0 (striping) in that it stripes the data across multiple drives. Where it's different is that if it's a 3-drive array, 1/3 of each drive is set aside to hold the parity information. If it's a 4-drive array, 1/4 of each drive is set aside for the parity information, and so on. So if you have a RAID5 array of three 80 GB hard drives, it will look like one 160 GB hard drive (approximately) to the operating system. If one of those drives fails, you can replace it and rebuild all the data that was on the failed drive using the parity information on the other drives. This gives you significant performance gains, but still provides some extra safety for your data.

The important thing to remember is that nothing is ever fool-proof, so normal daily backups to tape or some other media still need to be performed. Besides, if you delete a file on a hard drive, and then later want it back, it might not be recoverable if other data has been written where that deleted file used to reside. Tape backups would be used to recover that file.

Oh, please keep in mind that the purpose of this document is to give you a good understanding of the technology, not to make you an expert. With that in mind, the explanations presented may not be 100% technically accurate, but will nonetheless give you a good working understanding of the technology.