

RAID Levels Explained

By Samara Lynn

ARTICLE DATE : March 27, 2014

pcmag.com

If you've ever looked into purchasing a [NAS device](#) or [server](#), particularly for a small business, you've no doubt come across the term "RAID." RAID stands for Redundant Array of Inexpensive (or sometimes "Independent") Disks. In general, a RAID-enabled system uses two or more hard disks to improve the performance or provide some level of fault tolerance for a machine—typically a NAS or server. Fault tolerance simply means providing a safety net for failed hardware by ensuring that the machine with the failed component, usually a hard drive, can still operate. Fault tolerance lessens interruptions in productivity, and it also decreases the chance of data loss.

The way in which you configure that fault tolerance depends on the RAID level you set up. RAID levels depend on how many disks you have in a storage device, how critical drive failover and recovery is to your data needs, and how important it is to maximize performance. A business will generally find it more urgent to keep data intact in case of hardware failure than, for example, a home user will. Different RAID levels represent different configurations aimed at providing different balances between performance optimization and data protection.

RAID Overview

RAID is traditionally implemented in businesses and organizations where disk fault tolerance and optimized performance are must-haves, not luxuries. Servers and NASes in business datacenters typically have a RAID controller—a piece of hardware that controls the array of disks. These systems feature multiple SSD or SATA drives, depending on the RAID configuration. Because of the increased storage demands of consumers, home NAS devices also support RAID. Home, prosumer, and small business NASes are increasingly shipping with two or more disk drive bays so that users can leverage the power of RAID just like an enterprise can.

Software RAID means you can setup RAID without need for a dedicated hardware RAID controller. The RAID capability is inherent in the operating system. Windows 8's Storage Spaces feature and Windows 7 (Pro and Ultimate editions) have built-in support for RAID. You can set up a single disk with two partitions: one to boot from and the other for data storage and have the data partition mirrored.

This type of RAID is available in other operating systems as well, including OS X Server, Linux, and Windows Servers. Since this type of RAID already comes as a feature in the OS, the price can't be beat. Software RAID can also comprise virtual RAID solutions offered by vendors such as Dot Hill to deliver powerful host-based virtual RAID adapters. That's a solution more tailored to enterprise networks, however.

Which RAID Is Right for Me?

As mentioned, there are several RAID levels, and the one you choose depends on whether you are using RAID for performance or fault tolerance (or both). It also matters whether you have hardware or software RAID, because software supports fewer levels than hardware-based RAID. In the case of hardware RAID, the type of controller you have matters, too. Different controllers support different levels of RAID and also dictate the kinds of disks you can use in an array: SAS, SATA or SSD.

Here's the rundown on popular RAID levels:

- **RAID 0** is used to boost a server's performance. It's also known as "disk striping." With RAID 0, data is

written across multiple disks. This means the work that the computer is doing is handled by multiple disks rather than just one, increasing performance because multiple drives are reading and writing data, improving disk I/O. A minimum of two disks is required. Both software and hardware RAID support RAID 0, as do most controllers. The downside is that there is no fault tolerance. If one disk fails, then that affects the entire array and the chances for data loss or corruption increases.

•**RAID 1** is a fault-tolerance configuration known as "disk mirroring." With RAID 1, data is copied seamlessly and simultaneously, from one disk to another, creating a replica, or mirror. If one disk gets fried, the other can keep working. It's the simplest way to implement fault tolerance and it's relatively low cost.

The downside is that RAID 1 causes a slight drag on performance. RAID 1 can be implemented through either software or hardware. A minimum of two disks is required for RAID 1 hardware implementations. With software RAID 1, instead of two physical disks, data can be mirrored between volumes on a single disk. One additional point to remember is that RAID 1 cuts total disk capacity in half: If a server with two 1 TB drives is configured with RAID 1, then total storage capacity will be 1 TB not 2TB.

•**RAID 5** is by far the most common RAID configuration for business servers and enterprise NAS devices. This RAID level provides better performance than mirroring as well as fault tolerance. With RAID 5, data and parity (which is additional data used for recovery) are striped across three or more disks. If a disk gets an error or starts to fail, data is recreated from this distributed data and parity block— seamlessly and automatically. Essentially, the system is still operational even when one disk kicks the bucket and until you can replace the failed drive. Another benefit of RAID 5 is that it allows many NAS and server drives to be "hot-swappable" meaning in case a drive in the array fails, that drive can be swapped with a new drive without shutting down the server or NAS and without having to interrupt users who may be accessing the server or NAS. It's a great solution for fault tolerance because as drives fail (and they eventually will), the data can be rebuilt to new disks as failing disks are replaced. The downside to RAID 5 is the performance hit to servers that perform a lot of write operations. For example, with RAID 5 on a server that has a database that many employees access in a workday, there could be noticeable lag.

•**RAID 6** is also used frequently in enterprises. It's identical to RAID 5, except it's an even more robust solution because it uses one more parity block than RAID 5. You can have two disks die and still have a system be operational.

•**RAID 10** is a combination of RAID 1 and 0 and is often denoted as RAID 1+0. It combines the mirroring of RAID 1 with the striping of RAID 0. It's the RAID level that gives the best performance, but it is also costly, requiring twice as many disks as other RAID levels, for a minimum of four. This is the RAID level ideal for highly utilized database servers or any server that's performing many write operations. RAID 10 can be implemented as hardware or software, but the general consensus is that many of the performance advantages are lost when you use software RAID 10.

Other RAID Levels There are other RAID levels: 2, 3, 4, 7, 0+1...but they are really variants of the main RAID configurations already mentioned, and they're used for specific cases. Here are some short descriptions of each:

•**RAID 2** is similar to RAID 5, but instead of disk striping using parity, striping occurs at the bit-level. RAID 2 is seldom deployed because costs to implement are usually prohibitive (a typical setup requires 10 disks) and gives poor performance with some disk I/O operations.

•**RAID 3** is also similar to RAID 5, except this solution requires a dedicated parity drive. RAID 3 is seldom used except in the most specialized database or processing environments, which can benefit from it.

- **RAID 4** is a configuration in which disk striping happens at the byte level, rather than at the bit-level as in RAID 3.

- **RAID 7** is a proprietary level of RAID owned by the now-defunct Storage Computer Corporation.

- **RAID 0+1** is often interchanged for RAID 10 (which is RAID 1+0), but the two are not same. RAID 0+1 is a mirrored array with segments that are RAID 0 arrays. It's implemented in specific infrastructures requiring high performance but not a high level of scalability.

For most small- to midsize-business purposes, RAID 0, 1, 5 and in some cases 10 suffice for good fault tolerance and performance. For most home users, RAID 5 may be overkill, but RAID 1 mirroring provides decent fault tolerance.

It's important to remember that RAID is not backup, nor does it replace a backup strategy—preferably an automated one. Backing up to a RAID device might well be a part of such a strategy. Owning a RAID-enabled device, which you use as your primary server or storage device, is not. RAID can be a great way to optimize NAS and server performance and quickly recover from hardware failure, but it's only part of an overall [disaster-recovery solution](#).

Copyright (c) 2015 Ziff Davis Inc. All Rights Reserved.