

To make informed disk array purchase decisions, it is important to look beyond a discussion of the different RAID levels and also understand the differences between software and hardware-based RAID implementations. Software-based arrays are typically the least expensive to implement. Conversely, the true cost of this implementation can increase dramatically if the demand they place on the server CPU necessitates an upgrade to maintain acceptable network performance.

Software & Hardware RAID Defined

Software-based RAID implementations are either operating system-based, or they are application programs that run on the server. Most operating system RAID implementations provide support for RAID-1, and to a lesser degree, some operating systems also provide support for RAID-5. All array operations and management functions are controlled by the array software running on the host CPU. Most hardware-based arrays are implemented directly on a host-based RAID adapter and tightly couples the array functions with the disk interface. Additionally, this design allows all of the array operations and management to be off-loaded from the host CPU and instead, be executed locally on an embedded-processor. Different hardware RAID adapters offer some combination of RAID configuration support options including RAID-0, RAID-1, RAID-3, and RAID-5.

Performance Differences in Hardware & Software RAID

Just like any other application, software-based arrays occupy host system memory, consume CPU cycles and are operating system dependent. By contending with other applications that are running concurrently for host CPU cycles and memory, software-based arrays degrade overall server performance. Also, unlike hardware-based arrays, the performance of a software-based array is directly dependent on server CPU performance and load.

Except for the array functionality, hardware-based RAID schemes have very little in common with software-based implementations. Since the host CPU can execute user applications while the array adapter's processor simultaneously executes the array functions, the result is true hardware multi-tasking. Hardware arrays also do not occupy any host system memory, nor are they operating system dependent. Hardware arrays are also highly fault tolerant. Since the array logic is based in hardware, software is NOT required to boot. Some software arrays, however, will fail to boot if the boot drive in the array fails. For example, an array implemented in software can only be functional when the array software has been read from the disks and is memory-resident. What happens if the server can't load the array software because the disk that contains the fault tolerant software has failed? Software-based implementations commonly require a separate boot drive, which is NOT included in the array.

Differences In RAID-5 Implementation

Disk I/O is the most expensive operation in terms of system resources that any operating system can perform. In a RAID-5 environment, each write causes the array subsystem to perform up to three additional I/Os, plus two parity calculations, to complete the original write. Overall server performance is bottlenecked since software-based RAID uses the host CPU for parity calculations and forces the operating system to perform all of the additional disk I/O that's generated by the 'single' RAID-5 write, in addition to handling all application and operating system functions.

In a hardware RAID-5 configuration, the operating system is only required to perform the original write. The hardware-based RAID adapter processes all of the additional RAID-5 I/O, thereby off-loading the host CPU and freeing it to more adequately handle user applications and operating system functions. Only a 'true' hardware RAID solution avoids the I/O bottleneck.

The bottom line of this is that hardware RAID, while initially more expensive than software-based RAID, is the best selection for overall system performance. It is also less expensive than software RAID when the true cost of implementation, including the cost of server upgrades that may be required due to increased load, are considered.