

Tagged Command Queuing (TCQ)

Data Transfer Optimization in WD Raptor™ EL74 Serial ATA Hard Drives

In servers and network attached storage, it is commonplace for many users to read and write to the same hard drive simultaneously. These random reads and writes can be performed much faster if they are intelligently ordered to read/write to/from the nearest disk sectors. Intelligently ordered (queued) commands avoid additional revolutions of the hard drive (see Figure 1 below) and greatly improve performance for random I/O workloads commonly found in multi-use systems such as servers and network attached storage.

A common comparison is to imagine an elevator proceeding to the floors in the order the buttons are pushed rather than proceeding to the nearest floor. For elevators carrying only one or two passengers (analogous to a non-random PC workload) this is not a big problem; but for elevators carrying lots of passengers (analogous to a random I/O server workload), the impact to the passengers is obvious.




Figure 1. Command Ordering in Tagged Command Queuing

Tagged Command Queuing

This same common sense is applied to hard drives and data access with Tagged Command Queuing (TCQ). The adapter adds tags to the individual commands. The hard drive reorders the commands based on seek distance and rotation. The drive has dedicated buffers that will accept up to 32 commands at one time. The drive can then internally change the order of processing the commands to optimize the seeks much like an elevator.

In a hard drive there is another component of mechanical delay related to the rotation of the disks and where the data is relative to the head. This rotational latency is also considered in command reordering to further optimize the performance of the drive. An example of this would be two outstanding commands that are similar seek distances from the current position of the heads. Based on the rotation of the media, the drive determines one of the seeks would arrive just after the data passed and the other would arrive just before the data sector arrives at the head. It would then choose the second option as the data would be available sooner and decrease the rotational delay.

Tagged Command Queuing Benefits

- Random IOPs increase with or without write cache.
- Designed to excel in multi-threaded environments with high random I/O workloads.
- The drive's firmware handles the ordering of overlapping/colliding queued commands; otherwise remedied by the controller.
- Performance benchmark data available to compare WD Raptor to SCSI.

Conclusion

Command queuing has been standard in the SCSI environment for many years. Serial ATA has created new opportunities for performance improvement at lower cost. By adding Tagged Command Queuing to SATA drives, Western Digital has closed the gap in performance compared to higher priced SCSI drives.