Improving SSD Performance: *Combating Write Amplification to Extend SSD Lifespan*

Overview

Storage technology is constantly evolving. One of the latest technologies is Solid State Drives (SSD) - which replaces traditional electro-mechanical parts (i.e. rotating disk platters and read/write heads) with flash memory.

Benefits of SSD storage include:

- Very fast random access times due to elimination of slow electro-mechanical components
- Low read latency times due to elimination of disk seek times
- Consistent read performance because physical location of data doesn't matter (there is no "fastest" part of the drive as in traditional disk drives)
- File fragmentation has a negligible effect due to the elimination of electro-mechanical components (seeking)

Due to the nature of the flash memory and how data is currently written, SSD write performance degrades over time.

The Inherent Flaw with Flash Memory

Every write to a solid state drive results in the undesirable action known as “write amplification,” which shortens the life of the drive.

The flash memory of a solid state drive must be erased before it is rewritten, so when you try to write to the SSD, user data and metadata must be rewritten more than once to accomplish the intended write.

The multiplied number of writes and the bandwidth consumed by write amplification have two negative consequences:

1. Decrease the lifespan of the SSD
2. Reduce random write performance

But there is something you can do to slow down, if not reverse, the effects of write amplification.

To improve SSD performance and extend the life of your SSD you must limit writes to the drive. There is an easy way to accomplish this but let's dive into the issue of “why” first.

SSD Performance Factors

SSD performance depends on the following factors:

- **Write Endurance**[1]**:** The number of write cycles to any block of flash is limited. The maximum number of write cycles (endurance) is dependent on type of flash memory (MLC vs. SLC) and varies from 10,000 write cycles in older SSD drives to 1,000,000 write cycles or more with today's modern SSD drives.

- **Write Amplification**[2]**:** Write amplification is native to all NAND flash memory. Just as with traditional disk drives with NAND flash memory, data is laid down in blocks. However, block sizes on an SSD are fixed - meaning even a small 4k chunk of data write can take up a 512k block of space, depending on the NAND flash memory being used. When any portion of the data on the drive is changed, a block must first be marked for deletion in preparation of accommodating the new data (read/modify/write). The amount of space required for each new write can vary. The write amplification factor on many consumer SSDs is anywhere from 15 to 20. That means for every 1MB of data written to the drive, 15MB to 20MBs of space is actually needed[3]. *For example, a read/modify/write algorithm in an SSD controller will take a block about to be written to, retrieve any data already in it, mark the block for deletion, redistribute the old data, then lay down the new data in the old block.*

[1] http://en.wikipedia.org/wiki/Flash_memory#Write_endurance
[2] http://en.wikipedia.org/wiki/Write_amplification
[3] Knut Grimsrud, a director of storage architecture in Intel's research and development laboratory.

To maintain SSD write performance, SSD manufacturers implement one or more of the following techniques:

- **Wear Leveling**[4]**:** The SSD controller keeps track of how many erase cycles have been performed on each flash block and dynamically remaps logical to physical blocks to spread out the wear over all the cells in the drive. This means that no one portion wears out faster than another - prolonging the life of the SSD.

- **Over Provisioning:** Over provisioning provides extra memory capacity (which the user can't access). The SSD controller uses these "extra" cells to more easily create pre-erased blocks - ready to be used in the virtual pool.

- **TRIM**[5]**:** TRIM allows the SSD controller to remove data from deleted cells so that the next write won't have to move, erase and then write. This allows an SSD to maintain write performance for a longer period of time. In order for TRIM to be effective, it has to be implemented in the SSD itself as well as in the Windows operating system. TRIM has been included with all new operating system releases since Windows 7 and Windows Server 2008 R2.

[4] http://en.wikipedia.org/wiki/Wear_leveling
[5] http://en.wikipedia.org/wiki/TRIM_(SSD_command)

How to Combat Write Amplification

Because SSDs read data from flash memory, file access and “read” times are not a problem with SSDs...but writing to the disk is another issue.

SSD users should be concerned about the fragmentation of free space. If free space is scattered across the SSD between full blocks of data and trapped within partially full blocks of data, the more places the SSD must look to in order to write to the disk, and the less efficient write operations become.

You can combat the effects of write amplification by keeping free space consolidated on the SSD. **Write amplification actually decreases** when running TRIM operations to free up disk space, so you want to use the TRIM command to wipe clean unused disk space trapped in partially full blocks of data.

But how do you keep TRIM operations running as efficiently as possible?

Is It OK to Defrag an SSD?

While there has been much discussion around whether or not to defrag an SSD - the consensus is in:

Do not defragment your SSD.

Since there is no mechanical seek time on an SSD, traditional file-based defragmentation really doesn't provide any performance benefit.

In fact, you can do more harm than good by performing a defrag on an SSD, as it would actually create additional writes to the drive -- *unlike a hard disk drive, any write operation to SSD storage requires not one step, but two: an erase followed by the actual write* -- so SSD defragmentation should be avoided.

With SSD storage, the whole idea is to decrease the number of writes/updates to the SSD, so you want to be sure that any sort of optimization pass performed on the SSD does as little "shuffling" of files and data as possible.

But Yu Hsuan Lee of Apacer Technology, a company that produces industrial SSD solutions, wrote an article at *RTC Magazine*[6] that discusses and provides some benchmarks showing that optimizing an SSD drive improves performance and extends an SSD's life span:

> ***"Since the erase/write speed is slow compared to read, a write multiplication due to free space fragmentation can slow down I/O time severely."***
>
> ***"This means a well-designed defrag algorithm can extend an SSD's life span."***

Intelligent SSD Optimization

High free space fragmentation is a strong indicator that a high instance of untrimmed -- or partially full -- blocks exists on an SSD. Free space consolidation eliminates free space fragmentation and consolidates partially full blocks of data. This results in more efficient TRIM operations and faster write performance, reducing write amplification.

The new requirement for managing SSDs is a disk optimizer that identifies which drive is an SSD and which is a traditional hard drive and then performs the appropriate actions for each drive.

[6] http://rtcmagazine.com/articles/view/101053

PerfectDisk® and Solid State Drives

PerfectDisk's SSD Optimize feature, specifically designed for SSDs, automatically eliminates free space fragmentation and consolidates fragmented free space wherever the largest section of contiguous free space exists, whether at the beginning, middle or end of the drive.

While PerfectDisk is known for its efficient defragmentation and fragmentation prevention on traditional hard drives, its **SSD Optimize** feature **entirely avoids file defragmentation on SSDs**, focusing solely on the consolidation of free space. As mentioned above, file fragmentation does not inhibit SSD read performance, so running a traditional defrag would provide no benefit to the SSD.

Raxco Software draws upon its pioneering work with free space consolidation in PerfectDisk to provide **SSD Optimize**, which **focuses exclusively on what matters for performance in SSDs:**

- Consolidates free space on the drive **without** performing a traditional file defrag
- Identifies where the largest section of free space is located and consolidates free space in that location -- regardless of whether it is at the beginning, middle, or end of the disk

PerfectDisk detects SSD hardware and defaults to the SSD Optimize setting for SSDs. Running SSD Optimize on your solid state drive automatically results in more efficient SSD TRIM operations, preventing the multiple writes caused by write amplification before they occur, leaving you with faster, more efficient writes to the disk and a longer-lasting drive life.

By working with the beneficial properties of solid state drive technology, PerfectDisk's SSD Optimize is able to maintain SSD performance over the long term without causing additional wear on the disk.

The automatic SSD optimization method SSD Optimize is included as a standard feature in PerfectDisk Professional and PerfectDisk Server.

About Raxco

Raxco Software is a Microsoft Gold ISV and a VMware Technical Alliance Partner member. For over 35 years Raxco has developed award-winning software solutions and performance utilities that improve performance on Windows servers and workstations, simplify administration and help our customers work more efficiently.