

Wear-Leveling and Life Span

What is flash drive wear leveling?

Flash wear leveling is a technique used to help prolong the life of Flash drives.

Data can be written to an address on a flash memory IC a finite number of times. Even though the allowable number of writes is large (typically 10,000 or 100,000), if you write to the same spot over and over, it is conceivable that the flash device would wear out at that location. So, flash drives use wear leveling to make sure this does not happen.

How does wear leveling work?

Wear leveling works to distribute data evenly across each memory block of the entire Flash drive. This process decreases the total wear on the drive, thereby increasing the lifetime of the drive.

In order to understand how wear leveling works, one needs to understand the basics of how memory is stored and addressed on a flash drive. A couple terms to be familiar with are LBA and PBA.

LBA (Logical Block Address): The LBA is the address that the operating system uses to read or write a block of data on the flash drive.

PBA (Physical Block Address): The PBA is the fixed, physical address of a block of data on the flash drive.

The Flash Drive Controller resides on the flash drive, and provides the lookup table to map which PBA contains the data assigned to a given LBA. This relationship is shown in Figure 1. Since a lookup table is used, the data for a given LBA can be stored at nearly any PBA on the drive.




Figure 1

How does the “leveling” happen?

As we have discussed, the controller uses a lookup table to map the LBA to the PBA; it is a lot like the index of a book. So, the data can physically be moved, and all one has to do is change the index, and the data can still be found with ease.

Since one can move the data around with little or no penalty, the controller constantly does just that. Updated or new data is written to the first available free block with the least number of writes. The block that contains old data is erased in the background and then marked as a free block. This block rotating technique ensures even wear of memory blocks across the Flash drive. The wear leveling process is transparent to the Operating System.

The following is a very simplified example, illustrated in figure 2. At the start, we have PBA #1 through #4 containing data, and PBA #5 and #6 are “free”, or empty. When the computer writes new data to LBA15, the controller writes the new data to PBA #5, instead of PBA #1, where LBA15 used to reside. PBA #1 is now “free”.

Now, we will write to LBA14. As with LBA15, the data is written to a new PBA, and PBA #3 is now free. Finally, we will write to LBA14 a second time. Once again, the LBA is moved to a new physical address. Using a far more complex version of this activity, the writes to the flash drive are made fairly constant across the device.




Figure 2

What is static wear leveling?

This technique wear-levels over both dynamic and static data areas (see figure 3). For example, an 8GB Flash stores 3GB of MP3 music files. The entire 8GB capacity can be used for wear leveling.

Reassigning static data is a more complex task than dynamic data because it requires multiple operations to safely move static data around. As a result, this implementation may impact the overall write performance.

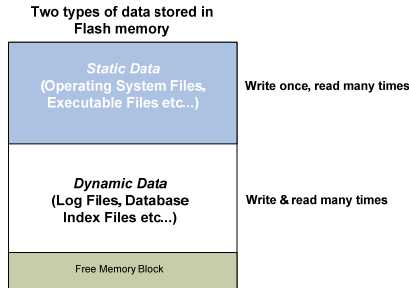


Figure 3

What is *dynamic* wear leveling?

This technique only wear-levels over dynamic data areas (see figure 3). For example, an 8GB Flash stores 3GB of MP3 music files but only 5GB remaining capacity is free for wear leveling.

The dynamic data areas are rotated in the round-robin fashion from a pool of free blocks. Dynamic wear leveling has a shorter life expectancy compared to static wear leveling because only dynamic data areas can be used to rotate data.

Static Wear Leveling vs. Dynamic Wear Leveling : Tradeoffs

Item	Static	Dynamic
Endurance	Very long life expectancy	Long life expectancy
Performance	Slower	Faster
Design Complexity	More complex	Less complex