WP21-0720-01

SSD Endurance Calculation Details

Introduction

SSD endurance is used for gauging the longevity or life of an SSD. Engineers often use this parameter when comparing SSDs. Its unit is measured in terabytes, often displayed as TBW for terabytes written. Sometimes, it is referred to as total bytes written, which is technically more correct. Regardless, of the terms, it is necessary to keep track of the units when calculating SSD endurance.

Endurance basically is the number of bytes that can be written to the SSD before it becomes a read-only device. This paper goes into details how endurance is calculated without having to inquire the SSD manufacturers.

Endurance

The equation below shows the formula for endurance found in almost all storage literature and is often shown as TBW.

Equation 1:

$$TBW = \frac{\text{Raw NAND Capacity x NAND PE Cycles}}{\text{WAF}}$$

NAND PE Cycles rating is provided by the SSD manufacturer, which is guaranteed by the NAND manufacturer and therefore is fixed. The two unknowns left of the equation is to determine the Raw NAND Capacity and the Write Amplification Factor (WAF).

Raw NAND Capacity

To determine the raw NAND capacity of the SSD, let's look at the architecture of the NAND itself:

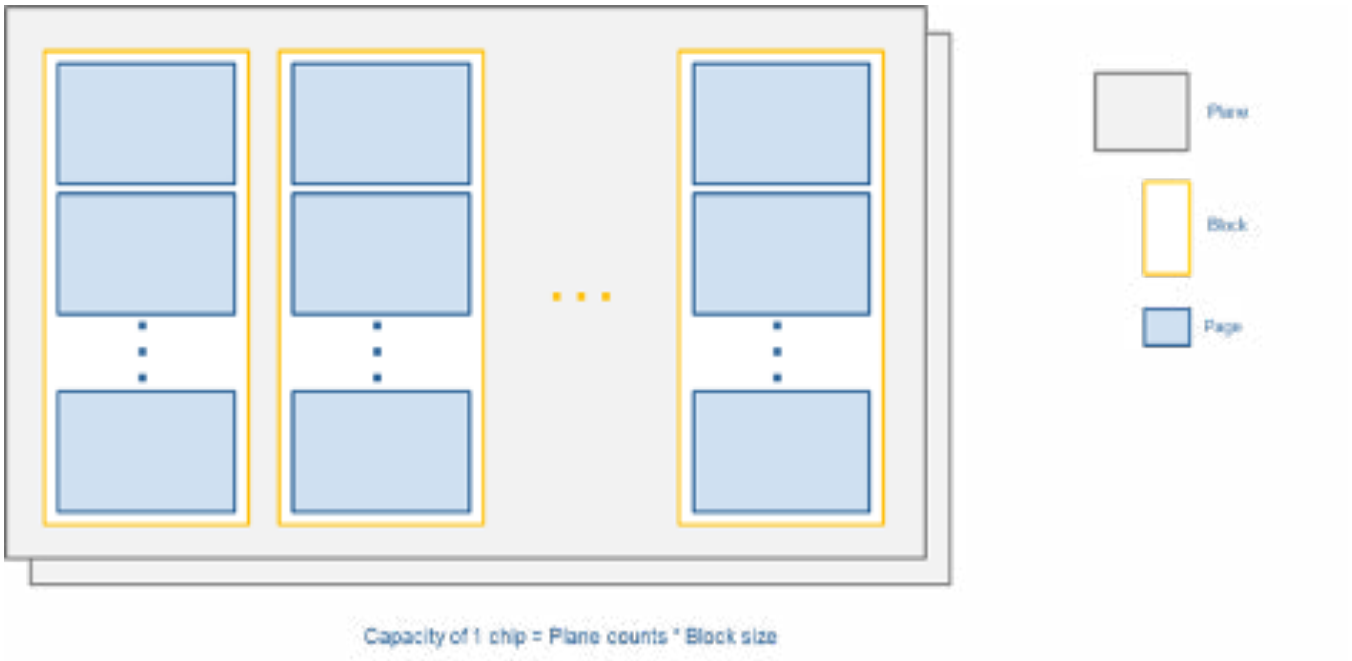

SSD NAND consists of planes, blocks, and pages. Numbers of pages together form a block. One plane consists of several blocks, and within the NAND chips, there are several planes. The capacity of the NAND will then be determined by the number of chips used.

Mathematically:

Block size = # of pages in a block * page size
Plane size = # of blocks in a page * block size
Capacity of 1 NAND chip = # of planes in NAND chip * plane size
Raw capacity of the SSD = # of chip * capacity of each chip

This approach is too cumbersome and most of the information is not available to the end user. Therefore, it would be impractical, and users generally resort to using typical values such as 32GB, 64GB, 128GB, 512GB, etc. This could prove erroneous especially if the SSD was overprovisioned to a non-standard value. For example, a 50GB SSD could be configured from either a 64GB or 128GB capacity. There is a better way, but first, let's determine how to find the WAF.

Write Amplification Factor

Equation 2 below describes WAF:

Equation 2

$$WAF = \frac{\text{NAND Writes}}{\text{Host Writes}}$$

Though it is not the intent to discuss various workloads and WAF in details, keep in mind that the same identical SSD will have a different WAF if each SSD was running with a different workload such as file sizes, amount of randomness, etc. Write amplification results from the NAND being written to more than what the host commanded.

SSDs with poor or higher WAF results in lower amount of data that can be written to the SSD. Therefore, by having a lower WAF, the SSD of the same capacity can have more data written to it compared to the one with a higher WAF. This is obvious with equation 1 referenced again:

Equation 1:

$$TBW = \frac{\text{Raw NAND Capacity x NAND PE Cycles}}{\text{WAF}}$$

As mentioned, the WAF is specific to one type of workload. As the workload changes so does the WAF. Therefore, whenever your system goes through an update or alters the software that changes how data is written to the SSD, it is likely that the workload will change.

One other important note to remember is that, depending on the capacity of the SSD, sufficient evaluation time should be taken to obtain a more accurate WAF. For SSDs with capacities up to 240GB, 20-40 testing hours is sufficient, while 120-200 hours is recommended for larger capacity SSDs. Allowing the SSD to run for at least 10% of its PE cycles will ensure that the SSD has gone through mechanisms that affects the WAF such as wear leveling, trim, garbage collection, and read disturb error management.

Now that we have established what WAF is, how exactly does it get measured? Within the controller, there is a SMART attribute. The SMART keeps records of the average erase counts for each block and the LBA written counts to the SSD.

NAND Writes

To determine the value of NAND writes, take the difference of erase count, and multiply it by the raw capacity in gigabytes (GB). For simplicity, we can let the average erase counts to be 0 at time zero t_0.

Difference of Average Erase Counts = average erase counts at t_1 – average erase counts at t_0
Since we let the average erase counts at t_0 = 0
Difference of average erase counts = Δ average erase counts

Equation 3

$$\text{NAND Writes} = \Delta \text{ average erase counts x Raw NAND Capacity (GB)}$$

Host Writes

Host writes are recorded by the numbers of LBA written or “LBA Written Count” as reported in the SMART attribute. Similarly, to determine the difference of the LBA written count, assume that the initial count is 0 at t_0.

Thus,

Equation 4

$$\text{Difference of LBA written count} = \Delta \text{ LBA Written Count}$$
$$\text{Host writes} = \Delta \text{ LBA Written Count} \times 32 \text{ MB}$$
Here we multiply the LBA by 32MB since each LBA is 32MB

Substituting equations 3 and 4 into equation 2 yields

$$WAF = \frac{\text{NAND Write}}{\text{Host Writes}}$$

$$= \frac{\Delta \text{ Erase Count} \times \text{Raw NAND Capacity (in GB)}}{\Delta \text{ LBA Written Count} \times 32\text{MB}}$$

Thus,

$$WAF = \frac{\Delta \text{ Erase Count} \times \text{Raw NAND Capacity (in GB)}*1024\text{MB/GB})}{\Delta \text{ LBA Written Count} \times 32\text{MB}}$$

Equation 5

$$= \frac{\Delta \text{ Erase Count} \times \text{Raw NAND Capacity} \times 32}{\Delta \text{ LBA Written Count}}$$

Substituting equation 5 into equation 1 gives us:

$$TBW = \frac{\text{Raw NAND Capacity} \times \text{NAND PE Cycles}}{WAF}$$

$$= \frac{\text{Raw NAND Capacity} \times \text{NAND PE Cycles}}{\frac{\Delta \text{ Erase Count}*\text{Raw NAND Capacity} \times 32}{\Delta \text{ LBA Written Count}}}$$

Therefore,

Equation 6

$$TBW = \frac{\text{NAND PE Cycles} \times \Delta \text{ LBA Written Count}}{\Delta \text{ Erase Count} \times 32}$$

Equation 6 does not require the knowledge of the estimated WAF and Raw NAND Capacity of the SSD. After running the SSD for the recommended time length, both the “LBA Written Counts” and “Erase Counts” can be obtained from the SMART attribute.

Again, notice that the equation calls for the “differences” of the two, respectively. Therefore, it is imperative that at the start of the test, the initial values are recorded. The differences would be between the initial recorded value and the value at the end of the test, respectively.

Consequently, it is important that when a system goes through an update, it is best to assume that the workload will change, resulting in a different WAF from the prior workload.

For example, a user decided to update his system, but he wanted to know how much the Percent Remaining Life is left in the SSD after running it for two years. After checking the SMART attribute, the Percent Remaining Life is at 50%. It would be a mistake for him to think that the SSD will last another two years going forward. At best, it may last longer, but the downside is that it might last a lot less depending on how the WAF will be affected by the new workload.

An Even Simpler Solution

Virtium offers an even more convenient way to estimate the endurance of the SSD. Virtium's vtView software incorporates the information from the SMART attributes and generates graphical views of the data. Users can visualize how the workload and any software updates impacts the SSD endurance.

vtView offers users with essential history data to better understand any possible issues that may have occurred during the debugging process.

Click here to learn more and about vtView, and how it can simplify your tasks.

Conclusion

While using equation 1 to determine the endurance of an SSD is sufficient, it requires users to inquire information such as WAF and Raw NAND Capacity from the SSD manufacturer. Furthermore, the published WAF value is based on a standardized test workload that is most likely different from that of the end users.

Equation 6, which utilizes the recorded values by the SMART attribute is the better and more precise method to estimate the endurance of the SSD. The WAF calculated by the controller is directly from the actual workload of the user.

Sufficient time should be taken to evaluate the SSD for a more accurate WAF measurement. A 960GB SSD may have a different WAF from the same SSD family that is 240GB if both were measured at 40 hours. It is recommended that the SSD runs for at least 10% of the PE cycles.

An even more convenient method to simplify your task in estimating the SSD endurance can be done with Virtium vtView.

Virtium manufactures memory and storage solutions for the world's top industrial embedded OEMs. For two decades we have designed, built and supported our products in the USA - fortified by a network of global locations. Our world-class technology and unsurpassed support provide a superior customer experience that continuously results in better industrial embedded products for our increasingly interconnected world.

Virtium
Solid State Storage and Memory

30052 Tomas | Rancho Santa Margarita, CA 92688
Phone: 949-888-2444 | Fax: 949-888-2445
www.virtium.com