Options for High SSD Endurance

Introduction

When it comes to SSD endurance, SLC NAND is the definite winner. However, as new technologies and new processes become available as well as with the ever-increasing cost of the SLC NAND, some OEMs are considering the 3-D TLC NAND.

With proper overprovisioning (OP), 3-D TLC NAND can offer respectable endurance and substantial savings compared to the 2-D SLC.

Endurance Calculation

Unlike hard disk drives, solid state drives (SSD) have limited program/erase (PE) cycles. Most industrial grade SSDs have PE cycles ratings from 3K and up to 10K cycles for TLC. Depending on the capacity, and types of workload, an incredible number of terabytes can be written to it before the drive is retired.

The endurance of the SSD, often measured in terabytes written (TBW), depends on 3 factors: capacity, NAND PE cycles, and write amplification factors (WAF). Basically,

Equation 1:

$$TBW = \frac{\text{phycal capacity x NAND PE Cycles}}{\text{WAF}}$$

Where,

Equation 2:

$$WAF = \frac{\text{amount of data written to the NAND}}{\text{amount of data written by the host}}$$

Looking at the TBW equation, only write application factor (WA or WAF) ultimately affects the endurance since the SSD capacity and flash PE cycles could be considered constants if comparing similar SSDs. Write amplification factor are mostly affected by file size and the types of workload. In general, the WAF is basically around 1 for sequential workload. The mix workload is where WAF will vary. This paper compares the differences in endurance based on a JEDEC enterprise workload with various OP.

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).

Overprovisioning

Overprovisioning basically is the reserve space taken from the NAND's capacity during firmware programming. It often increases the performance and endurance of the SSD. The negative side of OP is loss of capacity, however.

Consider the data in Table 1 below; these products have the same physical BOM, just configured differently at various OP.

Table 1 OP vs. WAF

NAND Configuration	Raw Capacity	User Capacity	OP	WAF	NAND PE	TBW
(2x128GB) TLC	256 GB	240 GB	7%	7.95	3,000	97
(2x128GB) TLC	256 GB	200 GB	28%	2.56	3,000	300
(2x128GB) TLC	256 GB	100 GB	156%	1.31	3,000	590

The trend clearly shows that the higher the OP, the higher the endurance. If capacity is not an issue, configuring the 256GB NAND to 156% OP is the obvious choice.

Cost consideration

In most cases, engineers and OEMs generally have a fixed budget and often need to abide to the immediate cost savings. Therefore, if the 100GB capacity suffice their application needs, it would be more cost effective to use a lower capacity (2X 64GB: 128GB) instead of higher capacity (2x 128GB: 256GB) drive and configure it to 100GB. On the other hand, to appreciate the true value of the higher endurance option, we need to look at the dollars per TBW instead of dollars per GB. For example, the cost of 256GB raw capacity SSD is around $100 while the 128GB is about $70. To compare the $/TBW, the endurance of the 128GB drive will need to be evaluated.

The OP of 100GB from 128GB raw NAND capacity is 28%. According to the measured value at 28% OP, the expected WAF is about 2.56 as seen in table 1 above. Applying the endurance equation with WAF=2.56, the TBW for the (2x 64GB) 128 raw capacity drive is about 150GB, since it is half that of the 256GB. Refer to Table 2.

Table 2 TBW of 128GB vs 256GB raw NAND capacity

NAND Configuration	Raw Capacity	User Capacity	OP	WAF	NAND PE	TBW
(2x128GB) TLC	256 GB	200 GB	28%	2.56	3,000	300
(2x64GB) TLC: HE	128 GB	100 GB	28%	2.56	3,000	150
(2x128GB) TLC: XE	256 GB	100 GB	156%	1.31	3,000	590

For the sake of discussion, let us refer to the (2x 64GB) and (2x 128GB) products from the table above as HE-100 and XE-100 respectively for high and extra high endurance.

Comparing the endurance based on user capacity, XE-100 is 590 TBW compared to 150 TBW for HE-100. Cost-wise the XE-100 is $100/590TBW ($0.17/TBW) compared to $70/150TBW ($0.47/TBW) for the HE-100 SSD.

Now let us consider the drives writes per day (DWPD):

Equation 3:

$$\text{Endurance: DWPD (x years)} = \frac{\text{Endurance (TBW)}}{\text{Rated Capacity * x years * 365}}$$

DWPD (HE-100) = 1.37 for 3 years

DWPD (XE-100) = 5.39 for 3 years

The XE-100 product could provide nearly five times the DWPD over the HE-100 product. If your application requires 100GB capacity with 1.3 DWPD for three years, the HE-100 will easily fit the requirement at an immediate lower cost. However, if you decide to go with the XE-100, 590 TBW could last you 12 years. So instead of replacing the drive after three years,your system can run an additional nine years. The true savings in this case is invaluable. Many changes could take place year after year as products may become obsolete, price may change, etc.

Even Higher Endurance

One other option available to achieve an even significantly higher endurance and a boost in performance is to operate the SSD in pSLC mode. The impact, however, is a much larger capacity loss (2/3 of the raw NAND capacity when configuring a TLC).

Using the same NAND capacity of 256GB as pSLC, that would reduce to roughly 85GB. From that capacity some space would need to be reserved for the firmware program and to improve wear leveling, and ultimately improve the overall write amplification. A 70% OP is used in the following example.

Despite the lower raw capacity (85GB vs 256GB) and slightly higher WAF than the XE-100 product, the overall endurance of the pSLC product is significantly higher since its NAND PE has increased to 30K from 3K.

Table 3 Endurance comparison of high OP and pSLC products

NAND Configuration	Raw Capacity	User Capacity	OP	WAF	NAND PE	TBW
(2x128GB) TLC	256 GB	100 GB	156%	1.31	3,000	590
(2x64GB) TLC: pSLC	85.3 GB	50 GB	70%	1.58	30,000	1620

Performance-wise, there is about a 45% improvement in sequential writes and 5x faster random writes per the test conditions.

Table 4 Performance comparision of high OP and pSLC products

NAND Configuration	User Capacity	Throughout		IOPS	
		Sequential Read (MB/s)	Sequential Write (MB/s)	100% Random Read	100% Random Write
(2x128GB) TLC: XE	100 GB	540	300	67K	8K
(2x128GB) TLC: pSLC	50 GB	550	450	72K	45K

Conclusion

The calculations in this paper provides an arbitrary cost and reliability requirements. Actual required warranty periods and capacities ultimately depends on the engineers and OEMs to decide which option suites their needs.

Engineers should also consider the overall cost savings of the program rather than the immediate savings alone. When targeting longevity, the high OP and pSLC products may be a cheaper solution.

Virtium's sales representative is available to discuss and determine the best solutions in terms of cost, longevity, and performance for your needs.

Virtium
Solid State Storage and Memory

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

WP23-1220-01