DPACK: Drive Writes Per Day (DWPD)

Briefing on how to leverage DPACK's *Average Daily Writes* value to calculate SSD durability **2016**

Will my SSDs wear out?

The three biggest objections to implementing SSDs in an environment are cost, speed, and durability. Luckily, the industry is providing comfort in all three of these categories.

This briefing will touch on price as well as performance, but will focus on durability: specifically, the topic of DWPD or Drives Writes Per Day, which has become a standard for indicating wear or durability expectations of SSD drives.

Price and Capacity

NAND is the underlying technology you will find in USB or SSD "Flash" and it's being made available at a continuously reduced price. The drivers for this cost reduction are changes in the way NAND is manufactured. Two common practices are (1) increasing the number of bits per cell, resulting in MLC and TLC technology, and (2) what is referred to as 3D or V-NAND, which is the technique of vertically stacking those cells. TLC and 3D techniques are often combined to create today's very high-capacity and cost-effective SSD drives.

Performance

While SSD capacity and price benefited from these advancements in manufacturing, performance and durability were challenged by enterprise adoption. There is a major argument about write performance on high bit cell NAND (TLC) SSD due to the longer programming cycles of the TLC NAND. However, the SSD write performance is dominated by the SSD SoC (system on chip) and it's firmware which mitigates much of that concern.

Increased capacity has helped overcome this problem as well. Drives are rarely 100% full and firmware can take advantage of that fact to perform a process called *Garbage Collection,* which proactively prepares write space so that the preparation penalty is mitigated when a write occurs. Since SSDs don't share the seek time penalties of HDDs, using space anywhere on the drive is just as optimal as any other location. As a result, today's high capacity 3D TLC drives are surprisingly fast.

Also, a single digit percentage of companies actually need the IOPS that most marketing claims would have you believe and you're probably not one of them. If you want to prove that to yourself, just run DPACK and see. Higher capacity and more cost-effective drives bring the benefit of migrating a greater amount of production data on Flash to provide a consistent quality of service to IO across more data capacity.

DWPD or Drive Writes Per Day

The combination of most companies having much lower IO needs than led to believe, mixed with higher capacity drives, can dramatically change one's opinion on adopting 3D TLC drives as tier 1 capacity drives.

DWPD simply is the number of times you can completely overwrite the capacity of an SSD per day and stay within the manufactures' recommendations. In this next section, you are going to learn how to calculate it for your unique environment.

In This Briefing

- The 3 reasons SSDs are often avoided
- Understanding DWPD

© 2016 Dell

DWPD: Using the average daily writes value to calculate SSD endurance

All Flash technologies inherit one flaw which is that the process of writing to Flash gradually degrades the memory cells. Storage administrators must consider the write workloads of their applications before deploying SSD disk and caching products to ensure that the product lifetimes match their requirements.

The standard for measuring the endurance of a SSD disk is Drive Writes per Day (DWPD). DWPD is measured in terms of the total capacity of the disk. For example, a 100 GB SSD does one DWPD if the user writes 100 GBs in one day. The standard suggests that the disk will endure the estimated DWPD for at least 5 years.

DPACK can assist users by showing at each layer (disk, server, cluster disk, collector run, and project) the estimated average daily write.

Average Daily Write

To calculate the average daily write for any given set of IO records, simply add the write throughput (MB/sec) of all of the records. Also, sum up the duration of each record. This is done automatically for you in DPACK and results in a capacity value that will be written each day. The basic equation to understand Average Daily Writes would be:

$$AverageDailyWriteMB = \frac{SumAllWriteMBperSec}{SumAllRecordSeconds} * SecondsInOneDay$$

Therefore, for any SSD drive type one can use this capacity value in combination with the following equation to calculate the minimum number of drives needed to accommodate the daily write activity, including any backend IO operations:

$$MinimumNumberActiveDrives = \frac{(AverageDailyWriteMB * RaidPenalty)}{(DriveCapacity * DWPDRating)}$$

Note: In this equation *"RaidPenalty"* will need to be explained in further detail later in this document.

To utilize this equation, one must first know the drives DWPD rating. This rating is issued by the manufacturer and is generally available data associated with the specifications of that drive. For the purposes of demonstration, here are some generally acceptable DWPD ratings for various drive types.

Type	Drive Size	DWPD
SLC	400GB	10
SLC	800GB	10
MLC	480GB	3
MLC	1.9TB	3
TLC	960GB	1
TLC	3.8TB	1

Often times a higher DWPD is achieved by over-allocating capacity and can contribute to higher costs.

As an illustration of this, a 200GB SSD with a DWPD rating of 3 can often be equivalent to a 600GB SSD with a DWPD rating of 1.

To exceed either DWPD rating an equal amount of 600GB per day would need to be written.

Up Next… Understanding RAID's impact on calculating DWPD

© 2016 Dell

RAID: A quick primer on how RAID effects DWPD

RAID 10 is the simplest form of RAID to understand. With each write and additional copy is written to the other disk in the mirror. Therefore, the RAID penalty used will be 2. RAID 5 & 6 are more complicated and at first glance RAID penalties used in DWPD calculations might look at odds to commonly-held knowledge of "capacity efficiency ratios," but this can be easily understood with a simple diagram as they are related, but mutually exclusive capacity factors.

RAID 5: RAID Penalty is 2

The useable capacity efficiency ratio for a RAID 5 (4+1) is 80%. Simply put, 4 capacity disks and 1 parity disk is a 4/5 ratio. DWPD is calculated based on capacity of written data, but more importantly how it's written to disk. To illustrate we will need to introduce some common RAID terms. Each RAID set will consist of a *RAID Stripe Width* and *RAID Stripe Depth*.

RAID Stripe Width: Is the number of Drives the RAID stripe will span. In this case 4 disks + 1 parity disk.

RAID Stripe Depth: This term can have many names, but is the quantity of data that will be written to each disk before writing moves to the next disk. *This is a critical factor in understand the logic of estimating DWPD.*

The diagram below shows the best- and worst-case scenario of writing to this RAID stripe with a 64KB stripe depth.

Best Case Scenario:

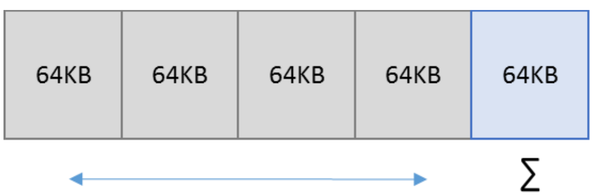

Systems may try to do write concatenation or coalescing to attempt to optimize the impact to disk. In this case a perfect 256KB is written. Each disk would receive an even allocation of 64KB. Parity would also be 64KB, but is only a 20% overhead of the 256KB written.

Worst Case Scenario:

However, most writes are small and often much smaller than the stripe depth. In this case let's assume only an even 64K of data has been written. This would only effect two disks in the RAID stripe: the disk where the 64KB was written and the re-written parity disk which is also 64KB, resulting in a 100% write overhead, even though usable capacity would remain 80% efficient.

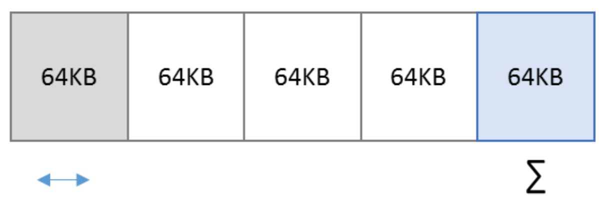

RAID 6: RAID Penalty is 3

An extra penalty has to be taken for dual parity RAID. In the same worst-case scenario, the 64KB would be written and only effect one disk; however, now two parity disks would need to be recalculated and re-written. Therefore, the worst-case scenario for RAID 6 will be a penalty of 3x (64K of data + 128K of parity).

Summary

Since DWPD is a factor of data capacity written and we can assume that an SSD will try to optimize writes by finding a new pre-prepared portion of the drive vs. overwriting the same data space. We can create an extremely safe estimate of the wear to any given drive with any known Average Daily Write capacity.

The calculations on the following page will reflect a 100% worst-case scenario estimate, so any optimization to writes provided by the system will only make these estimates safer.

Up Next… How to apply these calculations to every day scenarios!

© 2016 Dell

DWPD: Applying *Average Daily Write* to estimate endurance

There are two ways to approach the DWPD value depending on what you are trying to achieve: calculating the Minimum Number of Active Disks needed or Calculating the Estimated Life Span of a particular amount of SSDs with a known workload.

Minimum Number of Active Drives

Introduced earlier, this method will help you estimate if one drive or a set number of drives will be needed to stay within the recommended DWPD rating based on an observed workload demand in a DPACK project.

Example Data:

Average Daily Writes: 3.5TBs
Destination RAID Set: RAID 10
SSD being evaluated: TLC 3.8TB SSD

$$1.8\, Drives = \frac{3584GB * 2}{3891.2GB * 1}$$

Average Daily Writes: 3.5TBs
Destination RAID Set: RAID 5-5 (4+1)
SSD being evaluated: TLC 3.8TB SSD

$$1.8\, Drives = \frac{3584GB * 2}{3891.2GB * 1}$$

Average Daily Writes: 3.5TBs
Destination RAID Set: RAID 6-6 (4+2)
SSD being evaluated: TLC 3.8TB SSD

$$2.8\, Drives = \frac{3584GBs * 3}{3891.2GBs * 1}$$

What is demonstrated here is that even with the extreme amount of writes at 3.5TBs a day, and with the RAID penalty included, the high capacity of the drives results in the required minimum number of disks to be 3 or less for all the configurations.

Smaller Drives and higher DWPD

To illustrate the relationship between disk capacity and DWPD Rating, the next example will use a much smaller drive with only 400GBs of capacity, but with a higher DWPD Rating of 10.

Average Daily Writes: 3.5TBs
Destination RAID Set: RAID 6-6 (4+2)
SSD being evaluated: SLC 400GB SSD

$$2.6\, Drives = \frac{3584GBs * 3}{400GBs * 10}$$

The end result is that the minimum number of drives to accommodate the write workload is still 3. However, the TLC configuration would have roughly 10TBs of raw capacity while the SLC would have 1200GBs.

Realized DWPD Factor

If you plan to use above or below the recommended number of minimum drives you accelerate or decelerate the wear estimate. Calculating this is as simple as dividing the recommended number of drives by the number you will use.

$$.27\, Realized\, DWPD = \frac{3\, Drives\, Recommended}{11\, Drives\, Configured}$$

This configuration is over-provisioned; therefore, the wear durability of these drives should exceed expectations.

$$1.67\, Realized\, DWPD = \frac{5\, Drives\, Recommended}{3\, Drives\, Configured}$$

This configuration does not meet the recommendation. Therefore, there is an accelerated wear on the drives.

Up Next… How to apply these calculations to estimate drive lifespan

© 2016 Dell

DWPD: Applying *Average Daily Write* to estimate drive life

To calculate the *Estimated Life Span* of a known number of drives against a known *Average Daily Write* capacity, one would reverse the calculations and use the Realized DWPD Factor.

Estimated Life Span

Most RAID sets, especially in a storage array, generally include 4-12 drives for a minimum configuration. The result is that calculating the life expectancy of a set of SSDs can often demonstrate a humorous estimated number of years the drives can be in service. However, in an improperly- or under-provisioned system, these estimates could help understand system refresh increments to avoid unexpected downtime.

Using our two Realized DWPD Factors from the previous page, one can demonstrate the effect of each scenario with this formula.

$$EstimatedLifeYrs = \frac{ManufacturerEstimate}{RealizedDWPD}$$

Example Data:

Manufacturer Estimated Years: 5
Realized DWPD: .27

$$18.5\ Estimated\ Years = \frac{5}{.27}$$

Manufacturer Estimated Years: 5
Realized DWPD: 1.67

$$2.9\ Estimated\ Years = \frac{5}{1.67}$$

Final Thoughts

For now, IOPS as a sizing measurement has been largely commoditized by the generally availability of SSDs at ever-decreasing prices and increasing capacities. The largest benefit of migrating to SSDs is getting more of your data to a consistently higher service level for all disk activity.

However, there has still been some hesitancy around the unknown as it pertains to speed and durability when it comes to a company's specific workloads and the innovative technologies SSD manufacturers are using to push the boundaries of capacity.

DPACK can measure the uniqueness of your environment and with this information you can find a comfort level around drive life expectancy no matter how unique your workload or what drive you choose to implement.

As a note, this document uses an Average Daily Write of 3.5TBs per day, which by all admissions would be far beyond the average corporation's demand in almost any vertical. So if your Average Daily Writes are below 3.5TBs per day, your wear expectations would exceed the estimates in this document.

If you want to know for sure, start today at https://DPACK2.Dell.com/register/DWPD for a complementary account.

DPACK is a vendor and platform agnostic standard method for getting the performance facts from your environment donated to the community by Dell, Inc.

© 2016 Dell

Understanding more about DPACK

The DPACK team is happy to conduct team or territory level training with employees, partners, or customers.

Our contact information can be found on the right in the blue column!

The DPACK site is located at https://DPACK2.Dell.com

The DPACK support site is located at https://DPACKSupport.dell.com or by emailing support@dpack.zendesk.com

The DPACK support site also has a vast library of other insightful attributes of DPACK and can be located here:

https://dpacksupport.dell.com/forums

Contact Us

The DPACK team can be contacted at the following address:

DPACKsupport.dell.com

General Manager:
Sam Kirchoff

NA:
Scott DesBles
Mike Bachman

EMEA:
Uwe Wiest

LatAM:
Alan Rabinovich

APJ:
Mike Bachman

DPACK: Know Your Load!

© 2016 Dell