Precision Workstation Storage Classification

By: Alex Shows, Dell Client Performance Engineering

This paper provides an overview of the method used for classifying disparate storage devices in order to differentiate high performance storage options for Dell Precision workstations.

Why workstation storage performance matters

Workstations are often equipped with some of the fastest components available to client systems. When it comes to storage, it's important to ensure that high-end professional applications aren't spending most of their time waiting for the storage subsystem. Workstation performance is only as a strong as its weakest link, and productivity lost waiting for reads and writes to the storage subsystem can significantly reduce the overall value of a workstation.

Storage classifications provide performance tiers

Storage devices, even those of the same capacity, mechanical form factor, physical electrical connection, and protocol may deliver significantly different performance levels. By classifying storage, performance tiers can be used to compare across storage devices which might otherwise lack clear specifications which denote performance. This is especially important for solid-state drives (SSDs), which can have dramatically different storage performance depending their internal architecture.

By classifying storage devices using the results of measured performance, Dell provides premium performance for professional workstation applications and the engineers that need it the most.

"Without measuring drive model performance in some systematic way, it can be difficult to differentiate SSD models."

Before SSD technology became mainstream, hard drive performance could often be ascertained by the so-called "back of box" specifications. That is, one could flip over the proverbial box and look at the spec sheet to read important performance factors like rotational speed, cache size, and average seek time. A detailed hard drive specification might even include other interesting factors like peak sustained throughput values and the number of platters inside the drive. While these specifications may not correlate directly with an individual application's performance, they do provide some basis for comparison and contrast across different drives in the rotational drive ecosystem.

Enter SSD technology: SSDs may have descriptive information in their product sheets, but it has become increasingly difficult to correlate SSD interface, flash memory type, and flash controller model to application performance. SSDs provide significant advantages over hard disk drives, but without measuring drive model performance in some systematic way, it can be difficult to differentiate SSD models based on their "back of box" specifications.

Measurement Methods

Performance measurements are essential to differentiating SSDs. But what should be measured? Reads and writes, both random and sequential, to be sure. Unfortunately however, if we were to answer this question by looking at some of the more popular client storage benchmarks, we might draw the conclusion that strictly peak performance should be measured. That is, in what ways can I measure the highest possible values for bandwidth or IO operations per second (IOPs).

Benchmarking Storage for Drive Classification

PCMark Storage

- Variety of consumer and office productivity workloads
- Measures using recorded traces of real applications which are played back
- Provides both an asymptotic score and an average bandwidth measurement

SPECwpc

- Targeted workstation applications, separated by vertical or industry
- Scale is based on a reference platform including an SSD
- Difficult to optimize as scoring is inversely related to average drive latency rather than peak throughput

IOmeter

- Robust scripting mechanism enables targeted measurements under tightly controlled access specifications
- Provides insights into storage performance scaling, for example with block size and queue depth
- Enables targeting of performance metrics that yield the most user benefit

These peak values and the methods used to measure them may or may not correlate to one's preferred professional application. This is especially true of benchmarks which measure peak performance under ideal application and environmental conditions such as high queue depth (the number of outstanding IOs in the storage pipeline).

In order to reflect a wide variety of workstation applications and usage models, the storage classification method we use incorporates multiple different types of measurements, to reflect both application-based as well as targeted synthetic performance.

Cluster Analysis

The results of performance measurements across all the various SSD drives are used as inputs into a cluster analysis. A number of different clustering algorithms were applied, including centroid-based and density-based, and the results of those were compared for fitness.

It became immediately clear that, likely due to the advantages of the PCI Express bus, NVMe SSDs easily formed a distinct grouping separate from SATA SSDs. This supported our earlier assumptions, and we continued to split the clusters within each of these interfaces.

SATA SSDs showed two clusters: a "mainstream" cluster centered on the average performing triple-level-cell (TLC) flash module, and a "performance" cluster centered on the average performing multi-level-cell (MLC) flash (see Figure 2). Here again, this division substantiated our earlier assumptions that MLC-based SSDs outperform TLC-based SSDs. However, in reviewing the SATA SSD clusters, significant overlap was present. That is, the lowest-performing drives in the high performing cluster might be the same or slightly lower performing in some measurements than the highest performing drive in the mainstream cluster.

NVMe SSDs also showed two clusters, and these were much more distinct than their SATA counterparts: the "mainstream" cluster was clearly occupied by TLC-based SSDs, while the "performance" cluster included the MLC and even a few SLC-based SSDs such as high-performance add-in cards (see Figure 1). Ultimately, we chose a centroid-based clustering algorithm to classify individual storage devices. Even though this created a bit of overlap between SATA SSDs with MLC flash and NVMe SSDs with TLC flash, it enabled us to simplify the classification system by deriving minimum performance metrics the high-performance classes.

Simplified Classification System

In order to provide clear guidance to SSD technology providers, it was important for Dell to establish some minimum performance guidance for the "performance" clusters. It was in this step that the storage classes came to be. Class 20 drives denoted mainstream SATA SSDs. These are the typical SSDs found in many client platforms and are suitable for workstation customers who want quiet, fast, and reliable drives but don't necessarily need the highest performance.

Class 30 drives represent the highest performing SATA SSDs. Initially when we introduced the classification system, there were quite many models of Class 30 SSD, but due to a number of factors including manufacturing costs and increasing performance of TLC-based drives, there are very few Class 30 drives today. Workstation users desiring higher performance than Class 20 should look to NVMe devices.

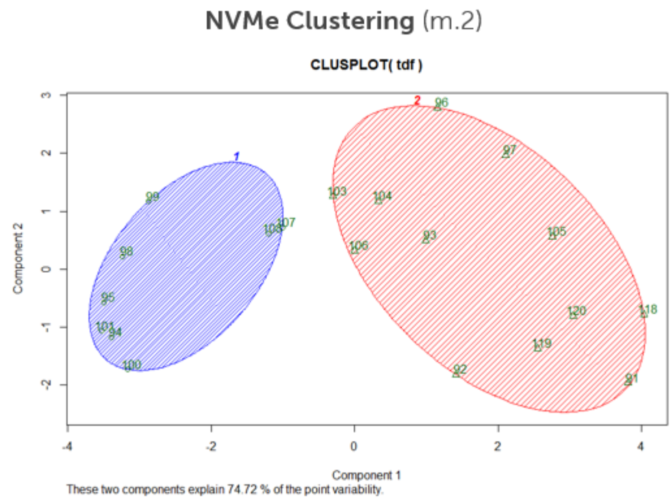

Class 50 is the **blue** shaded region

Class 40 is the **red** shaded region

Figure 1 *- Clustering of NVMe SSDs*

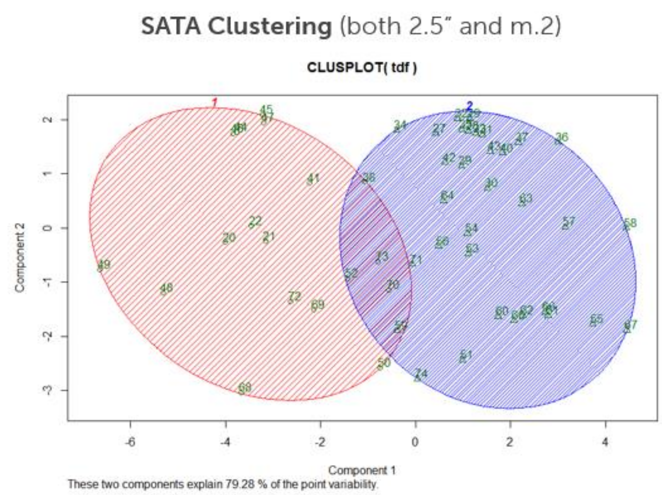

Class 20 is the **red** shaded region

Class 30 is the **blue** shaded region

Figure 2 *- Clustering of SATA SSDs*

Storage Classifications

Class 10

- Rotational disk drives
- Typically high capacity and low cost
- Typically SATA

Class 20

- SATA mainstream SSDs
- Significant performance increase over Class 10
- Comprised primarily of TLC-based SSDs

Class 30

- SATA performance SSDs
- Performance increase over Class 20
- Comprised primarily of MLC-based SSDs
- Fewer and fewer models available in this class

Class 40

- PCIe mainstream SSDs
- Comprised primarily of TLC-based SSDs
- Significant number of different SSD models available in this class
- Performance variation across different models

Class 50

- PCIe performance SSDs
- Significant performance increase over Class 40
- Comprised of MLC-based SSDs as well as newer flash-based technologies such as 3D Xpoint

Class 40 drives are predominantly mainstream NVMe devices that incorporate TLC flash. There are a significant number of different models available in Class 40, which tends to lead to performance variation as newer generations replace older generations. Class 50 drives are the highest performing SSDs offered on Precision workstations. They provide some significant performance gains over Class 40. Many are based on MLC flash, while some incorporate newer, faster technologies such as 3D Xpoint.

Class 30 Performance

In order to be classified as a high performance SATA SSD, a specific drive model (at a specific capacity point) must meet or exceed 8 out of 11 of the requirements listed in the following table:

Metric	Requirement
PCMark 8 Score	4900
PCMark 8 Bandwidth	220 MB/s
SPECwpc 2.0 Media & Entertainment	0.99
SPECwpc 2.0 Product Development	1.04
SPECwpc 2.0 Life Sciences	0.94
SPECwpc 2.0 Energy	1.11
SPECwpc 2.0 General Operations	1.08
IOMeter Random Reads 4K	39K IOPs
IOMeter Random Writes 4K	22K IOPs
IOMeter Sequential Reads 8K	247 MB/s
IOMeter Sequential Writes 8K	280 MB/s

Class 50 Performance

In order to be classified as the highest performing SSD, regardless of device connection, a specific drive model (at a specific capacity point) must meet or exceed 8 out of 11 of the requirements listed in the following table:

Metric	Requirement
PCMark 8 Score	5055
PCMark 8 Bandwidth	447 MB/s
SPECwpc 2.0 Media & Entertainment	3.17
SPECwpc 2.0 Product Development	2.98
SPECwpc 2.0 Life Sciences	2.65
SPECwpc 2.0 Energy	2.95
SPECwpc 2.0 General Operations	2.88
IOMeter Random Reads 4K	63K IOPs
IOMeter Random Writes 4K	44K IOPs
IOMeter Sequential Reads 8K	562 MB/s
IOMeter Sequential Writes 8K	765 MB/s

Classification Updates

Storage technology continues to advance, and the classification system is only useful if it advances with the technology. Every six months, we revisit the classification and look at variance within classes as well as outliers. As the centroids of the clusters move upward, this can sometimes affect the minimum required performance. This is why the number and variety of Class 30 drives is shrinking, while Class 40 continued to grow.

Conclusions

Solid-state drives provide significant performance advantages compared to rotational hard drives. This increased performance presents a challenge to identifying and distinguishing high performing SSDs from their mainstream equivalents. Dell’s solution to this challenge is a classification system, based on measured performance, to determine the minimum performance requirements for a drive to be classified as a performance SSD.