

QLC FLASH: THE LATEST ADVANCEMENT IN STORAGE



Keith Veigas

Associate Sales Engineer Analyst
Dell Technologies
Keithlennor22@gmail.com

Ritesh Agarwal

Associate Sales Engineer Analyst
Dell Technologies
R.agarwal@dell.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

Learn more at www.dell.com/certification

Table of Contents

Introduction	4
NAND flash	5
• Single-Level Cell	5
• Multi-Level Cell	5
• Triple-Level Cell	5
• Quad-Level Cell	5
Comparing the different types of NAND flash	6
Why QLC flash? Applications and Characteristics.....	8
Conclusion.....	11
Bibliography	12

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

Introduction

This article will address the advantages of Quad Level Cells (QLC) and try to understand under which workloads they will work best. To understand all of it, let's first learn a little about NAND (Not AND) in general.

Unlike other types of storage, programming data to a NAND cell requires erasing the old data in the cell before writing new data on it, referred to as the Program/Erase (or P/E) cycle.

How is the whole process carried out? It is done by sending electrons through an insulator back and forth by applying a voltage to it. The location and quantity of these electrons sent will determine the current flow between the voltage threshold (source and sink), which in turn will determine how much data (Binary: 0s and 1s) has been written/programmed in the NAND cell.

However, this procedure may have a drawback in certain situations. When the number of electrons will be more than what the insulator can take, it will have difficulty keeping the data in the correct location that we wanted to program it on. This will lead to one of the two scenarios; has the data been programmed to the correct location, or has it migrated somewhere else. To minimize this drawback and such scenarios, we have different types of NAND flash.

NAND flash

NAND comes in two types: Raw and Managed. This article will focus on the raw types of NAND.

Raw NAND comes in four subtypes:

- Single-Level Cell (SLC)
- Multi-Level Cell (MLC)
- Triple-Level Cell (TLC)
- Quad-Level Cell (QLC)

Raw NAND may need external management (by using an external host controller), but it is the most cost-effective (cost/GB) NAND flash in the market today.

- **Single-Level Cell:** As the name implies, SLC has 1 bit per cell, and has the highest performance amongst all other types of NAND flash. It is designed for high density and mission-critical workloads, where cost is not the major factor; rather, the major requirement is high performance and high durability/reliability.

When we store 1 bit, we need not keep a close eye on the location of the electrons, thus migration will not be a big concern. Since we are storing either a 0 or a 1, our only concern will be to determine if the voltage is flowing or not.

- **Multi-Level Cell:** MLC has 2 bits per cell and has a higher cell density compared to SLCs. Plus, it has a good level of performance and endurance.

Determining voltage threshold with MLCs is a bit complicated as it stores 2 bits per cell. Hence, the migration of electrons also has a greater impact here as compared to SLCs.

- **Triple-Level Cell:** Often considered the most common NAND flash used by a consumer today, TLC has 3 bits per cell and has a very high cell density as compared to both SLCs and MLCs. However, it has lower performance and endurance as compared to the other two.

The most appropriate workload for TLC will be for consumers' mass storage applications, for example, USB drives, which have very high-cost sensitivity.

- **Quad-Level Cell:** The most recently introduced in the NAND flash family, QLCs store 4 bits per cell, making it the highest cell density amongst all the other NAND flash types. However, the trend of higher the density, lower the performance continues here, and QLCs are the least performing cells and have the least endurance and reliability as well.

QLC is best used in enterprise applications with very high-cost sensitivity. In regard to migration of electrons, QLCs have the most significant impact since they store 4 bits of data per cell, which can be arranged in 16 different combinations.




Figure 1: Raw NAND subtypes

Comparing the different types of NAND flash

Figure 2 depicts the cell densities of the four types of NAND flash. As the number of bits per cell increase, so does the cell densities. Hence, SLCs have the least cell density, while QLCs have the highest cell density.

Inversely, as cell densities increase, performance, endurance, and reliability decrease. In this case, SLCs have the highest performance, endurance, and reliability amongst all NAND flash, whereas QLCs have the least. The below diagram represents the decrease in endurance threshold as we move from lower to higher cell densities.

Note that by levels, we mean the number of possible combinations in which each NAND flash can store data. It can be calculated by doing 2^n , where 2 denotes the binary digits 0 and 1, while 'n' denotes the number of bits that each cell can store data in.

Thus for SLCs, it will be 2^1 , i.e. 2. Similarly, for MLCs, TLCs, and QLCs, it will be 2^2 , 2^3 , and 2^4 . i.e. 4, 8, and 16 levels, respectively.

Understand that we can have a similar interpretation of decrease in reliability and performance as with endurance from SLCs to QLCs as shown in the diagram below.

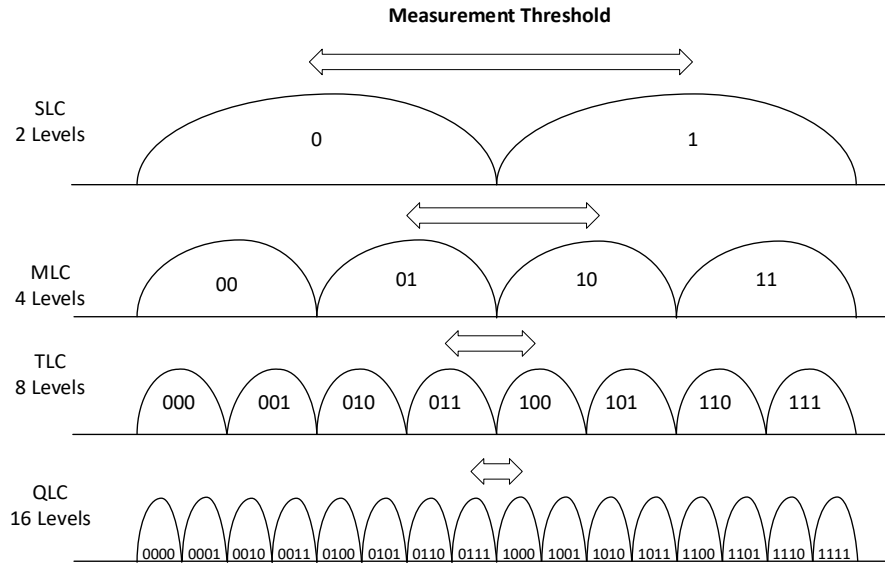


Figure 2: Cell Density for different types of NAND Flash

With an understanding that as the number of bits per cell increases, the cell densities increase, it's easy to see how this influences the program/erase (P/E) cycle per cell as well. In essence, moving from SLCs to QLCs, the flash becomes more read-centric and lesser write-centric. Where we can estimate around 100K P/E cycles for SLCs, as we move further, the number of cycles decreases significantly. For MLCs, the rough estimation is 10K P/E cycles, for TLCs it is 3K, and for QLCs it is only 1K P/E cycles per cell.

The next big question is cost. The prices of these flash devices are inversely proportional to their cell densities, i.e. the higher the cell density, the lower the price. This means SLCs, the first member of the NAND flash memory, are the most expensive (\$/GB) compared to MLCs, TLCs, and QLCs.




Figure 3: Value Proposition for different types of NAND Flash

Table 1 summarizes major characteristics of each of the NAND subtypes.

Metrics	SLC	MLC	TLC	QLC
Bits per cell	1	2	3	4
Cell Density	Least	Higher than SLC	Higher than SLC and MLC	Highest
Performance	Highest	Less than SLC	Less than SLC and MLC	Least
Reliability/Endurance	Highest	Less than SLC	Less than SLC and MLC	Least
Maximum P/E cycles	100K	10K	3K	1K
Cost (\$/GB)	Highest	Less than SLC	Less than SLC and MLC	Least

Table 1: Comparison of NAND subtype characteristics

Why QLC flash? Applications and Characteristics

For many years, read-intensive workloads have been primarily handled with legacy storage devices like hard-disk drives (HDD) since they are affordable for storing vast amounts of data. QLC (Quad Level Cell) NAND flash has become the new trend in the market mainly because it can store four bits in each cell as compared to three or fewer in earlier versions. This denser storage has also made flash more economical in current times, reducing the gap between HDD and solid-state drives (SSD) in terms of pricing.

In a modern data center, having lower rack space is important and QLC's denser storage facilitates immense savings in setting up the data center infrastructure. Also, having fast read-focused storage for various workloads helps bring critical data closer to the CPUs.

Applications

Many applications and workloads in the industry are more focused on reading than writing. QLC helps make the process of having this read-intensive workload on flash possible, thereby overcoming constraints with legacy storage options.

The biggest challenge in analyzing the most suitable workload for QLC is being aware of how much of your data and applications require for write operations. This limitation exists since QLC has lower write endurance compared to other Flash types.

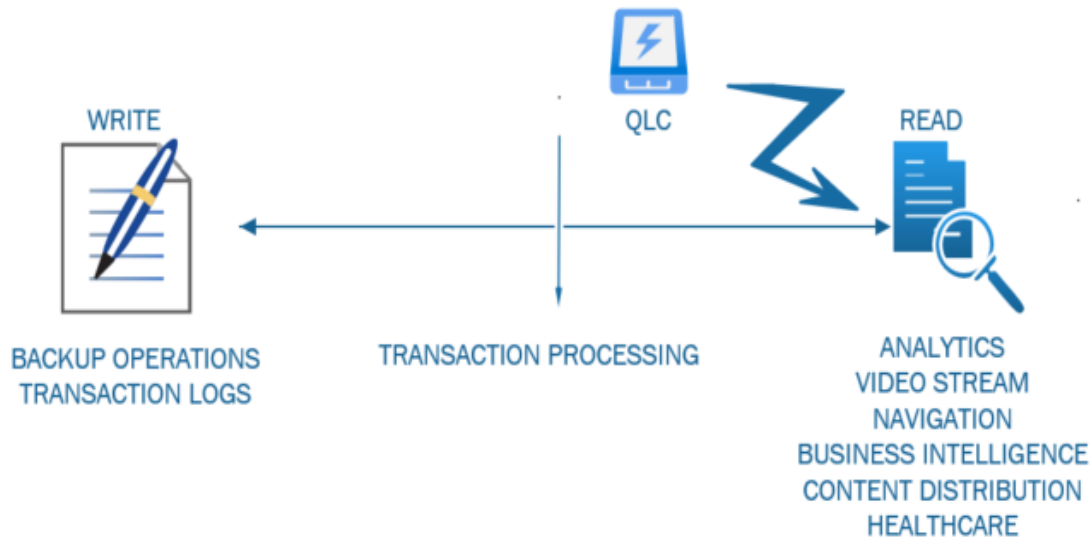


Figure 4: QLC-Flash Workloads & Use-cases

WORKLOAD	APPLICATION OF QLC
Machine Learning	QLC speed help read immense data sets quickly to analyze data.
Video Streaming	Can handle multiple requests to simultaneously stream data to more than one user at a given point in time.
Analytics	QLC provides a high performing backend for storage that is quick to access
Databases	In databases like NoSQL, tasks such as classifying content or tagging users are done seamlessly and efficiently.
Artificial Intelligence	AI algorithms require high-speed search to identify anomalies or patterns in huge data sets. QLC provides both large and fast storage for reads.
Business Intelligence	For fast mining of large datasets, QLC's can run deeper queries which enable more specific and simplified analytics resulting in better insights.
Active archival	Large archival data can be transferred quickly whenever needed from the storage to the processing unit.

Table 2: QLC applications to various workloads

Latency of data going to CPU

Having data closest to the CPU helps in processing it at maximum efficiency, lowest latency, and quick transfer of values. Unlike a traditional HDD, QLC flash has no moving parts ensuring that data always remains close to the CPU and available at all times. The higher latency on traditional HDD is mainly due to the mechanical operation of the drive, i.e. operations like random access require the head to scan the entire platter.

Have 33% more capacity per cell

SLC flash technology was introduced earliest and has remained the most expensive while at the same time being the fastest. After SLC, MLC took over where two bits stored in each cell was twice that of SLC. SLC made flash slightly affordable.

The 3rd revision, TLC is currently the most prominent technology, with 3 bits per cell, delivering a 50% increase compared to the previous generation. Introduction of TLC helped Flash become a mainstream product in many use cases due to lower costs and high-performance abilities.

Now, the emergence of QLC adds an extra bit per cell which contributes to a 33% increase compared to TLC, making flash even more affordable and a great replacement option for traditional storage.

Reduced Rack Space for large data set

Since more data can be stored per bit in a much smaller form factor, QLC solutions lead to incredible savings in rack space in a data center compared to having traditional forms of storage such as HDDs.

Conclusion

Considering current market trends, we expect NAND flash technology to be a major player in the coming days by producing even cheaper, faster, and more reliable products. As QLC becomes mainstream, traditional storage like HDDs will soon be replaced. Although QLCs offer less performance and reliability compared to all other types of NAND flash, their ability to be used with SLCs to further boost their performance and achieve respectable figures enables them to remain compatible with most workloads.

QLC makes having an all-flash data center at an economical price a reality. Since most of the workloads are read-focused, QLC becomes a perfect fit for having fast and cost-effective storage.

QLC is the top contender to be the next generation of flash storage. At the same time, we can predict greater advancements and further reductions in the cost of procuring a QLC drive.

Bibliography

Content and images

1. [Choosing the right NAND](#)
2. <https://docs.google.com/document/d/e/2PACX-1vT0PhEZQaoxmsa-SSpAVn1PMIA0I4Ty7BgXxSnichUPQf-xNq0wzn8oLSqmi75SRy0SY2UCWoHKU4oK/pub>
3. [Micron's solution to QLC](#)
4. <https://www.architecting.it/blog/qlc-nand/>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.