Micron® Quad-Level Cell Technology Advances New Affordable Capacity Standard

Latest-Generation NAND Technology Increases Pressure on Hard Disk Storage

Overview

We are constantly being challenged to build business services and solutions that provide maximum performance while fitting within tighter budgets. Whether buying compute, networking, or storage, each component should provide the right set of performance and features with optimized value. The traditional technology choice for low-cost capacity has been SATA hard disk drives (HDDs) due to their low cost per gigabyte, but this option has often resulted in compromises in performance and scalability. Micron believes that the solid state data center is inevitable. More and more, SSDs are becoming the mainstream storage medium of choice, but to date they have not been able to replace legacy hard drives in these business solutions that are heavily read-centric.

With the introduction of Micron's quad-level cell (QLC) NAND, the economics of SSDs move ever closer to HDDs for many applications. This technical brief compares QLC to its predecessor, triple-level cell (TLC) NAND technology and to SATA HDDs. When comparing QLC to TLC, we view QLC as a complementary technology to TLC and one that fills a gap between TLC-based SSDs and HDDs. The comparisons to HDDs show that in many use cases QLC is a viable option to consider as you move toward an all-flash data center. While QLC-based SSDs do not eliminate the cost-per-gigabyte advantage of HDDs, QLC does make a strong case for a massive consolidation of existing storage technology in favor of a more cost-effective one that offers better I/O, throughput, and latency performance — resulting in efficiencies that drive down the overall total cost of ownership (TCO) of large-scale, read-centric applications.

QLC Fast Facts[1]

- **33%** higher bit density than current TLC SSDs
- **32X** more write I/Os per watt than competitively situated SATA HDDs
- **2X** the throughput of competitively situated HDDs
- **23X** more data analyzed than competitively situated HDDs

The Performance Case for QLC as a Replacement for HDDs

What is QLC and why is it primed to challenge HDDs as an affordable capacity solution? QLC technology is the newest generation of NAND architecture, enabling each NAND cell in an SSD to store four bits of information, thus improving overall information density by 33% over current-generation TLC NAND packages. QLC is focused on read-centric workloads (90%+ read) that account for a large percentage of enterprise application use cases. Comparing QLC to SATA HDDs and current mainstream TLC SSDs shows that QLC-based SSDs can fill a very important gap between the two. QLC offers a unique opportunity to accelerate the transition from lower-performing HDDs to higher-performing SSDs at a more approachable price point than before. With read performance similar to that of TLC drives, QLC can support the consolidation requirements currently driving the industry, while providing over 450X the read IOPS (Figure 1) and 2X the read throughput of competitively situated HDDs when running the same workload (Figure 2).[1]

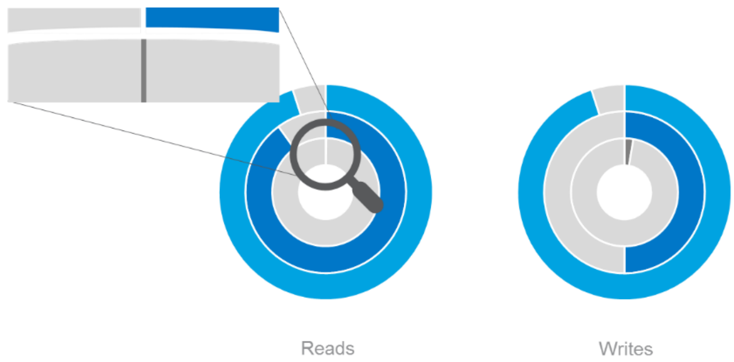

Figure 1: Performance (IOPS) Comparison of QLC SSDs vs. SATA HDDs and TLC SSDs

What does this mean for your data center? While HDDs continue to be the raw price-per-gigabyte leader, that is not the entire story. As illustrated, for both small-block (OLTP) and large-block (BI/DSS) type workloads, HDDs just cannot compete with TLC or QLC SSDs in terms of raw performance. This means you can manage the same data with fewer physical devices and may be able to do it faster, helping you to consolidate your application's hardware, software licensing, and power and cooling requirements for the same performance. This provides a massive equalizer in terms of cost per I/O or cost per MB/s.

For mixed I/O applications like OLTP or point-of-sale, lower-cost, high-performance SSDs can provide better value. For read-intensive applications, QLC can provide even better value and (as we'll see later) much better TCO.

For BI/DSS workloads that depend on large-block, sequential I/O, we can get more answers with half the equipment using QLC SSDs.[1] In these low-write environments, QLC even delivers better value than TLC SSDs, providing similar read performance at a lower price.

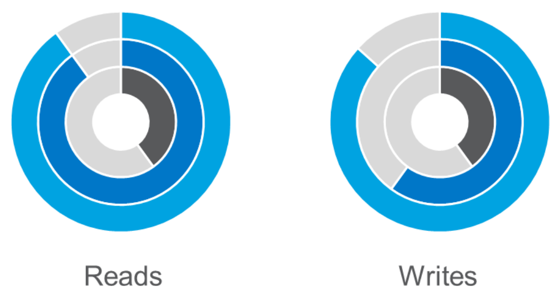

Figure 2: Throughput (MB/s) Comparison of QLC SSDs vs. SATA HDDs and TLC SSDs

From a latency perspective, SSD QoS (99.9+%) latency values are generally better than HDD *average* latency values, making them one of the most reactive storage solutions available. QLC-based SSDs will provide read latency similar to that of currently available TLC-based SSDs[1] but at a lower cost, making QLC a better overall option for read-intensive workloads.

All SSD performance results are based on actual testing using FIO with a queue depth of 32 for the stated read/write target. Micron SSD data sheets provide 100% read, 100% write and typically a 70/30 read/write for IOPS using 4KiB block size, as well as throughput in megabytes per second using 128KB block size at 100% read and 100% write only.

[1] Performance comparisons based on industry-agreed-upon performance and internally verified metrics for a typical enterprise performance-class 7.2K RPM SATA HDD compared to a Micron 5200 ECO SSD (TLC) (data sheet) and early results of Micron QLC-based 5210 ION SSD. Final 5210 ION specifications may be different from those used in this brief. All devices are 8TB devices.

The Efficiency Case for QLC as a Replacement for HDDs

When we consider the efficiency of what each of these devices does as part of a typical “day-in-the-life” of an application data drive, the story is just as compelling for QLC. Efficiency can be viewed in many ways, including power, space and cooling. These efficiency metrics contribute to the overall TCO advantage that SSDs, generally, and new QLC-based SSDs, specifically, offer.

Power efficiency is typically measured in terms of idle power consumed, maximum power consumed and number of operations (depending on the application) performed per watt of energy consumed. Since storage devices deliver business value by doing as much work as possible, this technical brief focuses on power efficiency while reading and writing data.

Unlike HDDs, SSDs do not have to physically move objects around as they read or write data. HDDs have to spin heavy platters and move read/write heads back and forth, which requires more electric power. Thus, SSDs will generally use less power than HDDs.

Maximum power metrics documented for HDDs range from 8 to11 watts, while maximum power consumption for SSDs is around 6 watts. While that may not be a major difference at the drive level, it is more enlightening when you evaluate the amount of actual work that can be accomplished with these power levels. Comparing the number of IOPS per watt that SSDs and HDDs can perform reveals a huge efficiency advantage for SSDs (Figure 3). While TLC SSDs will continue to be the raw performance-per-watt leader, QLC-based SSDs provide significantly higher power efficiency (32X) for writes than HDDs. In other words, it will take 32 HDDs to equal the power efficiency of a single QLC-based SSD!

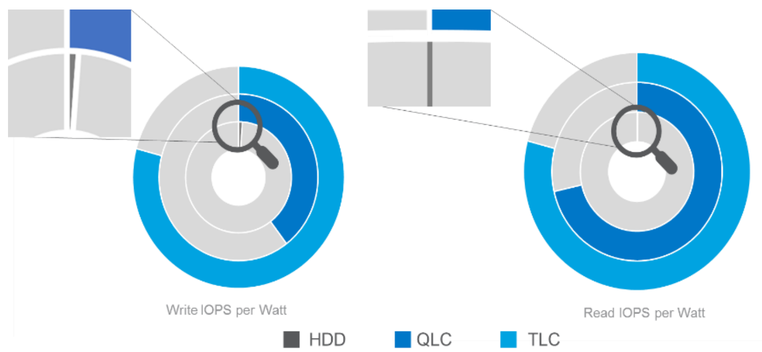

Figure 3: IOPS per Watt Advantages of SSDs

This efficiency opens up a massive opportunity for workload consolidation within the data center. Considering that QLC SSDs offer 2X the read throughput, 450X the read IOPS, and better power efficiency than HDDs, applications can be deployed on fewer servers with QLC SSDs versus HDDs in these solutions to get the same performance.

What if you need similar capacity? Even here, SSDs provide an advantage. Using our 8TB example, current HDDs come in a 3.5-inch-wide x 1-inch-high package, while all capacities of Micron enterprise SSDs come in smaller, 2.5-inch-wide packages. This means that a typical 2RU server chassis can support 2X the number of SSDs versus HDDs based on currently available designs. Using half the number of servers means less rack space and less data center capacity required, which also means less power and cooling. This results in a dramatic reduction in the TCO when using QLC SATA SSDs versus SATA SSDs for business intelligence and decision support systems (BI/DSS), video streaming and other highly read-centric workloads.

Assuming you are in the process of a five-year technology refresh, the average capacity of high-performance SATA HDDs five years ago was 2TB per drive. If you consider a refresh comparison of your existing three-server, 12 x 2TB HDD Cassandra solution versus a new deployment using 8TB QLC-based SATA SSDs, the charts below provide an estimated five-year TCO benefit for a performance-focused QLC-based SSD implementation of 77% over HDDs (Table 1), while the TCO benefit for an equivalent capacity-focused Cassandra solution shows a

QLC SDD advantage of 45% (Table 2).[2] Both QLC-based refresh strategies use fewer SSDs than the equivalent HDD-based implementations and still provide higher performance capacity to possibly support more transactions than the HDD-based Cassandra solution.

	Current	SSD	Difference
Number of Drives	36	2	94% Fewer Drives
Usable Capacity (GB)	72,000	15,360	79% Lower Capacity
Performance (IOPS)	6,516	8,178	26% Higher Performance
5yr TCO	$39,500	$9,042	77% Lower Cost

Table 1: Performance-Based 5-Year TCO Analysis

	Current	SSD	Difference
Number of Drives	36	10	72% Fewer Drives
Usable Capacity (GB)	72,000	76,800	7% Higher Capacity
Performance (IOPS)	6,516	40,890	528% Higher Performance
5yr TCO	$39,500	$21,776	45% Lower Cost

Table 2: Capacity-Based 5-Year TCO Analysis

The Endurance Case for QLC as a Replacement for HDDs[3]

As a storage technology, SSDs provide better lifetime endurance than HDDs, with lifetime endurance referring to the amount of data processed by a device. Currently available TLC SSDs can have a write endurance [measured in total bytes written (TBW)] of over 8 petabytes (PB), while enterprise-class SATA HDDs generally have a published “interface endurance” of around 2.5–3.5PB.[4]

So how is an HDD’s interface endurance different than an SSD’s TBW? HDD vendors typically publish limitations based on the maximum number of data bytes that can be passed through the drive interface in *either* direction (i.e., reads and writes), as well as a time-based warranty period, typically 3 or 5 years. Exceeding either metric can cause the warranty to expire. SSDs generally consider only the data written to the device, along with a time-based lifecycle and warranty period, typically 3 or 5 years. For this reason, all of an application’s I/O — both reads and writes — affect the lifetime of a SATA HDD, while only the data written by an application impacts the lifetime of an SSD. While QLC SSDs have a much lower TBW maximum than TLC drives, the lower TBW endurance doesn’t really matter for applications with high read-to-write ratios (90+% reads).

Consider an application that generates only 10% writes versus 90% reads and uses a 4KB block size, but is reading and writing data very rapidly, consuming the full IOPS capability of the drive. Using the I/O rates of HDDs, TLC SSDs, and QLC SSDs, all three options can manage this solution workload without exceeding their respective warranty limits for data throughput, even at these high I/O rates.

[2] Based on Micron-Forrester® Consulting MOVE2SSD TCO Tool metrics, using early 5210 ION performance and cost estimates for a 95%/5% workload hosted on a three-server Cassandra® cluster, running Yahoo! Cloud Server Benchmark workload B. Real-world results may differ from estimated values. Data used is subject to change.

[3] Endurance calculations are based on an industry-standard 5-year warranty for enterprise-class SATA drives used for comparison reasons only. Your warranty limitations may vary from those discussed in this brief.

[4] “Interface endurance” measures both read and write data that passes through the SATA interface between the host computer and the disk drive. HDD vendors typically place limits on the amount of data that drives can read and write compared to SSDs, which are based only on the amount of data written to, not read from, the SSD.

Drive Type	Write IOPS	Read IOPS	Total IOPS Used for Lifecycle Calculations	5Y Total Bytes Processed	5Y Total Bytes Maximum Allowable
HDD	21	189	210 (reads & writes)	135TB (reads & writes)	2750TB (reads & writes)
QLC	500 (estimated)	4500 (estimated)	500	322TB	450TB (estimated)
TLC	950	8550	950	613TB	8400TB

Table 3: 90%/10% Workload Impact on Endurance Warranty

For this example, the QLC SSD writes 23X as much data versus an HDD over the 5-year lifetime of both devices (Figure 4).

From a data throughput perspective (using large-block sequential data) with the same 90% read and 10% write I/O profile, QLC drives will take longer to reach their maximum data processing expectancy based on the SSD TBW maximum compared to the typical HDD interface endurance listed in an HDD vendor's warranty. Consider that HDDs can write data at a rate of 240 MB/s, while QLC SSDs can write data at a rate of approximately 340 MB/s. At these rates, the HDD would exhaust its expected interface data limit (reads and writes) in 139 days, while the QLC SSD would be expected to last for about 150 days.

As the read percentage increases above 90%, the QLC SSD advantage grows dramatically. In fact, at 100% reads with no data written to the device, the QLC SSD would never reach its maximum documented TBW limit in the SSD data sheet, but the HDD would exhaust its interface data limit (reads and writes) in the same 139 days (Figure 5).

An example of a 100% read workload would be a video streaming solution where, once the drive is filled with video files, it could be considered a 100% read workload from that point forward. QLC provides advantages in terms of both performance and endurance that maximize the TCO for this type of use case.

Figure 4: Comparison of GB Written Based on a 10% Write Workload Profile

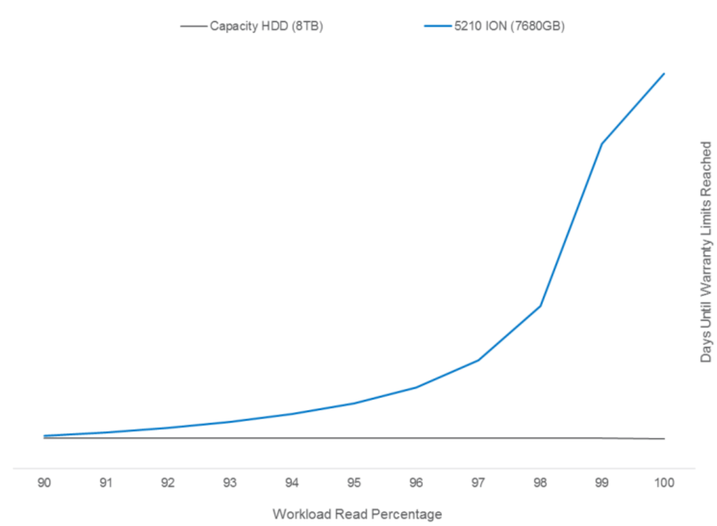

Figure 5: Days until Warranty Limits Are Reached for Large-Block, Sequential Workloads

Consider QLC for Your Next Read-Centric Application

Legacy HDDs have been a mainstay of low-cost, adequate performance, read-centric solutions due to their ubiquity and cost advantage over newer flash-based SSD devices. The introduction of QLC — the next generation of flash storage technology — provides much broader competition to HDDs with a lower-cost flash-based SSD option. While QLC SSDs cost more than HDDs on a per-gigabyte basis, QLC is the more cost-effective option for common read-centric workloads. The additional performance and efficiency tip the scales in favor of QLC SSDs over a five-year refresh cycle period of a typical enterprise solution.

The all-flash data center is closer to becoming a reality and, as technology continues to improve, it is all but inevitable. Now is the time to seriously consider QLC for read-centric workloads such as video processing, BI/DSS and archiving. Visit micron.com/QLC for more information about Micron's industry-first QLC SSDs.

Behind the Numbers

HDD Performance Assumptions

All performance comparisons use the following metrics for a "typical" enterprise 8TB, 7200 RPM, 3.5-inch SATA HDD:

Metric		Value
Throughput (MB/s)	Reads	240
	Writes	240
IO/Sec (IOPS)	Reads	198
	Writes	267
Latency (ms)	Average Read Latency	4.16
Reliability:	MTTF (Hours)	2 million
Power Efficiency:	Max Power (Watts)	8.6

IOPS per watt is calculated using the following formula:

Reads: $\frac{IOPSread}{Max\ Power}$

Writes: $\frac{IOPSwrite}{Max\ Power}$

Your specific models may offer different performance characteristics from those used in this brief.

SSD Performance Assumptions

All performance comparisons for TLC-based SSDs are based on published specifications for the Micron 7680GB 5200 ECO SSD as published in the technical data sheet available on micron.com.

All performance comparisons for QLC-based SSDs are based on early internal test results for the new Micron 7680GB 5210 ION SSD. Final product specifications may change from those used in this brief.

The following table summarizes the performance metrics used for each SSD type:

Metric		TLC	QLC
Throughput (MB/s)	Reads	540	540
	Writes	520	360
IO/Sec (IOPS)	Reads	95,000	90,000
	Writes	9,500	5,000
Latency (ms)	QoS Read Latency	1	.200
Reliability:	MTTF (Hours)	3 million	2 million
Power Efficiency:	Max Power (Watts)	6	6.3

5-Year TCO Estimation

TCO calculations are based on a hypothetical five-year refresh comparison of an existing three-node Cassandra NoSQL database solution, performing a typical 95% read-to-write workload such as adding metadata/tags to existing data, being replaced by a hypothetical QLC SSD-based offering with two possible decision criteria for the deployment:

1) Providing a new Cassandra solution that offers the same class of performance as the existing solution.
2) Providing a new Cassandra solution that offers the same capacity as the existing solution.

The tables below provide the assumptions used in the TCO calculations:

What type of HDD storage are you using now?

Note: Cost per drive is based on average public pricing.

Drive Type	HDD
Model	HDD_7200_RPM
Capacity	2000GB
Cost per drive	$138.33

How is your server storage deployed?

Chassis Size (in U)?	2
How many drives per chassis?	12
How many chassis like this?	3
What RAID level?	0
What % is in use? (HDD Only)	100%

What type of storage are you interested in?

Note: Cost per drive is based on average public pricing.

Drive Type	SSD
Class	Data Center
Vendor	Micron
Model	_5200_ECO
Capacity	7680GB
Cost per drive	$1,572.45

What Type of Application are you using?

Family	Database
Vendor	NoSQL_Cassandra
Workload	Adding metadata/tags to existing data
Custom Application?	☐
Transfer Size (KB)	0.5
Read/Write Ration (%)	100/0

Workload Details	
Transfer Size (KB)	128
Read/Write Ratio (%)	95/5

Software Licensing (per year)	
Current per-node licensing	$1,000
SSD-enabled per-node licensing	$0

Please enter your power costs and cooling factor

State	Use US Average
Custom Power Cost?	☐
$/KWh	$0.950
Cooling Factor	1.7

Power Cost	
US EIA $/KWh	$0.077

micron.com