KIOXIA

The State of the NAND Flash Industry and Achieving CapEx Efficiency

Scott Nelson | Executive Vice President & Chief Marketing Officer

If you're a design engineer looking to harness the power of cutting-edge flash memory technologies, understanding the current state of the NAND flash industry is crucial.

In this Q&A, I'll explore how recent developments are causing perspectives to shift on critical factors like CapEx. I'll also show how you can leverage these advancements to create more efficient, high-performance designs that meet the evolving demands of modern applications.

Q: Where does the NAND industry stand right now as it continues to emerge from the most severe downturn in its history?

A: Simply put – the NAND flash industry is 'back in black,' but the overall recovery continues.

The industry is still in active recovery mode and will be for a while – while ASPs continue to recover, we should take into consideration that the industry faced a $28B downturn, so we need to give it time for a full industry-wide recovery. But the timing for that recovery is very important, as that's going to help to determine future CapEx investments. Moving forward, the industry will be cautious with how it increases CapEx investments and is expected to take a more conservative approach.

As I mentioned in a previous conversation, we're now looking at technology migration in a completely different way. The CapEx intensity of leading-edge NAND technology nodes will slow migrations between technology generations, resulting in multiple NAND generations co-existing at any given time. This is a big shift from how we've handled migrations in the past, and we're calling it the 'New Normal.' I'll explain more a little later, but this is one way we will optimize CapEx efficiency in the future.

Q: What do you mean when you say CapEx efficiency?

A: Each successive generation of 3D NAND flash technology brings intensive, increasing CapEx expense along with it. Finding the balance between the relentless march of technological progress and the economic realities of skyrocketing costs is more important than ever. In the past, the industry would 'tick-tock' the entire volume/capacity from one NAND generation to the next. When you add the accelerated CapEx costs of leading-edge 3D equipment and the increasing footprint/ floor space required for it in the fab, it is no longer economical to move the entire capacity/ volume to the latest 3D generation.

What we've realized is - a good portion of the market doesn't need cutting-edge performance so there is no need to move everything to the latest generation. Therefore, maintaining multiple generations at a time and providing the latest generation to the portion of the market that actually needs leading-edge density and performance is the most efficient, sustainable way forward.

The new normal of supporting multiple 3D generations at once and making efficient utilization of existing/legacy capacity delivers a more competitive CapEx strategy. In addition, our use of groundbreaking CBA (CMOS directly Bonded to Array) technology, which separates the interface chip from the memory array, gives KIOXIA the technological agility needed to implement this strategy while delivering industry leading density and performance.

Q: KIOXIA is always pushing the boundaries of what's possible for flash technology – tell us about the newly introduced 2Tb QLC BiCS FLASHTM device. What market drivers are you addressing with this new device?

A: Our new 2Tb QLC device offers the industry's highest flash memory die capacity. The new device features CBA chip technology which we introduced in our latest generation of BiCS FLASH 3D flash memory. With this new device, we are able to offer the highest density memory solution for applications that demand the greatest levels of capacity and performance.

Our 2Tb QLC BiCS FLASH 3D flash memory addresses the pain points experienced by the cloud and data center markets including skyrocketing data growth and storage density, expanding footprints, escalating power usage, greater cooling requirements, increasing rack weights, etc. These segments need high density, power efficient storage solutions – and that's what we have delivered.

Over time, the demand for greater capacity and performance will continue to evolve, and with a stacked architecture - 16-die stacked in a single memory package - our latest QLC device achieves an industry-leading1 4TB of capacity.

Target Applications

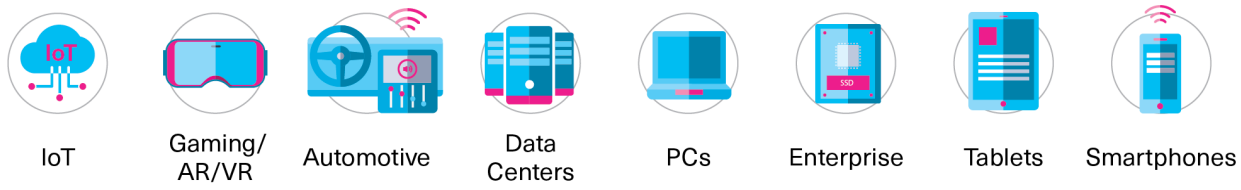

Q: You mention 4TB of capacity – it's difficult to understand how much capacity that is. Can you explain in practical terms?

A: It's pretty massive and hard to wrap your head around so let's take a look at a couple of everyday examples to put it into perspective.

Take podcasts, for example. A typical, 30-minute-long voice only podcast would be about 14.4MB in size. So, a 4TB device can store about 277,000 podcasts. It would take 8.3 million minutes to listen to them all (or 15.8 years if you listened 24 hours a day continuously).

Another example is simple text. The average English word is about 5 letters (6 if you count the space after each word). There are about 500 words in 11-point font on an 8.5 x 11 letter sized paper, so each page requires about 3kB. A 4TB device could store about 1.33 billion pages of text. If you laid 1.33 billion pages of text end to end, you would have enough to circle the earth 7 times. It would take many lifetimes to read this many pages, in fact, it would take 2,536 years of continuous reading to complete 4TB worth of text.

It's hard to comprehend that all of that data can fit in a single 13.50mm x 11.50mm package – about the size of your thumbnail.

Q: What can we expect to see in terms of solutions/applications?

A: With the growing demand and adoption of AI, data creation is projected to continue its rapid acceleration driving an ongoing appetite for high-performance storage solutions capable of addressing data center and enterprise challenges.

When it comes to data centers and enterprises, QLC-based SSDs will enable advancements in sustainability through improved power consumption, power efficiency, increased scaling density per watt enabling new levels of scalability in the data center.

QLC technology offers higher density, combined with performance and power efficiency, paving the way for innovative storage applications such as artificial intelligence, machine learning, edge computing, security/surveillance systems, and content delivery networks. Additionally, QLC is poised to play a significant role in client SSDs and mobile devices, presenting new opportunities for NAND technology.

> *The new normal of supporting multiple 3D generations at once and making efficient utilization of existing/legacy capacity delivers a more competitive CapEx strategy. In addition, our use of groundbreaking CBA (CMOS directly Bonded to Array) technology, which separates the interface chip from the memory array, gives KIOXIA the technological agility needed to implement this strategy while delivering industry leading density and performance.*

Scott Nelson, Executive Vice President & Chief Marketing Officer, KIOXIA America, Inc.

Notes:

The views and opinions expressed in this blog are those of the author(s) and do not necessarily reflect those of KIOXIA America, Inc.

1: As of July 2, 2024. Source: Kioxia Corporation

In every mention of a KIOXIA product: Product density is identified based on the density of memory chip(s) within the Product, not the amount of memory capacity available for data storage by the end user. Consumer-usable capacity will be less due to overhead data areas, formatting, bad blocks, and other constraints, and may also vary based on the host device and application. For details, please refer to applicable product specifications. The definition of 1KB = 2^10 bytes = 1,024 bytes. The definition of 1Gb = 2^30 bits = 1,073,741,824 bits. The definition of 1GB = 2^30 bytes = 1,073,741,824 bytes. 1Tb = 2^40 bits = 1,099,511,627,776 bits.