



Raj Hazra: Hello, and welcome to the Six Five Summit. My name is Raj Hazra. I'm a Senior Vice President and General Manager of the Compute and Networking business unit at Micron Technologies.

What I'm here today is to talk about the tremendous opportunity we all have, and, in particular, Micron has, in the data center. Why am I so excited? What we've considered as data centers, these big buildings with lots of computing equipment in them, are really becoming memory centers.

I have just a little bit of data and a lot of your own personal experiences to convince you that's true. If you look at the adoption of artificial intelligence in cloud today, you will see the nature of the impact of data and what it has done to data centers as we look through this data. In 2020, if you compare just over a span of five years from 2020 to 2025, the number of servers deployed in the cloud that deal with AI, some form of AI processing, whether it's training or inference has doubled in a mere five years. All the time, the total number of servers is going up, as well.

Data is really the raw material that AI shines on. As models get more sophisticated as with parameters going from hundreds of billions to trillions in just a matter of a year or year and a half, as usages become more sophisticated and integrated into critical elements of life, business, health, so has the data become the path to delivering the promise of AI. And where you have data, you have to have memory and storage. So it's no surprise that as more people need the benefit, want the benefit of AI, the number of AI optimized servers has grown.

Well, what that has done to memory and storage is compared to a compute-optimized server of yesteryear handling a number of traditional workloads. The AI workloads have created a different balance of compute and memory and storage. Storage and memory have grown seven X and six X respectively from what they would be if the server was not running an AI specific workload or not tuned for AI. So, these two things, the pervasive growth of AI and AI's inherent data dependency that drives memory and storage growth disproportionately than ever before, is what creates that tremendous opportunity for Micron in the global data center. As you see, in a matter of a decade, we see a 16% combined [inaudible] of the growth of memory and storage, particularly with storage growing a little bit faster, even, than memory in that timeframe. In essence, this is the time to be in memory and storage.

But, is it only because there's more of the same that will be needed? For folks that are in the business of the technology of creating in memory and storage, the second most exciting part of the world being data-centric and data centers becoming memory and storage centers is the amount of innovation that is going into memory and storage.

If you see even 10 years ago, there was a very simple, what I would call a system architecture to a system, that operated on multiple applications and workloads. We all remember the days when we talked about 640K of PCD Ram on a PC or a couple of megabytes on a simple server, but no longer is that the case. As these servers have become the host for a large number of diverse workloads, some compute-centric, many of them memory-centric, as they have dealt with data to do all of the analysis and inside generation, prediction, recommendations, create models, it has become very important to innovate in these memories to drive performance up, to drive cost down, to create new system architectures through the application of hierarchies of memory and storage.



No longer is it just direct attach and in-package hot memory for data that you're frequently using. You now have these layers of the hierarchy following it, including things like CXL attached memory, which is a new interface that allows far memory to be attached and expanding the server's capacity, as well as performance bandwidth, to deal with memory intensive workloads. And then, of course, storage, too, is going through this hierarchy [inaudible], as they call it, to create various levels of fast to cooler storage.

What this is being done for is, essentially, to answer the address of success of the data center. As data centers have to scale up to deal with the insatiable demand of this digital universe, handling more and more traffic, more and more data, more and more queries, more and more users, scale requires a tremendous amount of agility through software-defined efficiencies.

Inherently, computers have to become the very thing they do; be created and manage through software to provide performance through modularity and composability, driving to levels of efficiency, without which, frankly, we would not be able to have a practical implementation of a data center in the future. So, instead of racks of computer and memory and storage being statically defining the data center, what we see through innovations in memory hierarchies and storage hierarchies is the ability to create a data center on the fly.

By being able to associate memory pools with compute pools, move computation closer to where the data is versus moving data to where the computation has to be, an extremely important aspect of making these data centers energy efficient in addition to being performant. These completely changed the way that data centers will be architected, and this change will come about from the innovations of being able to create these different kinds of memory-based or memory-sensitive system architectures, all powered by innovations in the memory and storage media themselves.

This is why it's an exciting time, not only to be a beneficiary of what this can do for us, but to have the awesome privilege of being one that's involved in creating it, as well. When we talk about these exciting futures, the one thing we have to do is be able to span the gamut of needs of the different kinds of memory and storage. And this is where we at Micron are focusing on a complete data-centric portfolio, addressing the various markets that the data center is associated with, from the hyperscale cloud service providers, to the more traditional enterprise and government customers, to communication network providers and communication service providers, and even to the very fast and exciting growing intelligent edge.

The data center, really, is everywhere and/or everything is connected to the data center. It's running workloads of hundreds of difference of kinds, all needing memory and storage, all needing very optimized memory and storage. And so it's no longer a world of just DRAM or DDR. We have high bandwidth memory that we are starting to ship on the second generation and rapidly advancing to the third and fourth generations. We have the world's fastest graphics memory that we just recently announced, looking at low power DDR, not just in mobile phones anymore, but bringing them to not just the PC client space, but the data center, as well. The traditional DRAM advancing rapidly through DDR5, and then soon to DDR6, and on the storage side, going from TLC to QLC advancing NAND with the world's most complex and most advanced processes to help create these three-dimensional storage structures.



But. Just those things, those products, don't quite tell the story of what we need to do, what we have to do, and what we do every day here at Micron. We work very closely with our customers to understand their needs, to optimize to what they need. We standardize, we work through standards bodies, to ensure that the ecosystem is working together and ready to accelerate these technologies into reality. And then, of course, we enable the ecosystem to knock down barriers to adoption of new technologies, like we've done with the TEP program for DDR5.

All of this, fundamentally, is made possible by what I call the innovation substrate. We at Micron, today, are proud to say we are leading in silicon technology in both memory and storage with our 1-alpha technology for DRAM and our 176-Layer technology for NAND. In addition to the techno-leadership that brings tremendous advantages of cost reduction and capability to our products, we are also leaders in advanced packaging, bringing about heterogeneous integration of different kinds of logic and memory pieces, to create these advanced solutions, and also pioneering with significant investment in kind of the memories of the future.

Think of things that we can't do today, but we soon see the ability to do in the very near future by advancing, by attacking and solving some of the fundamental challenges of physics and chemistry, if you will, in creating these new structures for memory. All in all, it is an exciting time to be in the business, because what we do does indeed touch and change life for all, and that is the exciting part.

But we don't do it alone. We do it, and it's our commitment to do it, collectively with our partners in the ecosystem and through standards and other mechanisms for democratizing technology. Because our longstanding history, our proud history, has shown when we do that, we have the best way to bring the industry's best, including our own, in serving and delighting our customers. And that's really the story of the day, is coming to know who our customers are, serving and delighting them with technology innovations at a time when not only is what we do critical, it is absolutely essential for taking the world forward into the digital age. Thank you.