ReThink Data Center with Crossbar ReRAM technology

The Internet of Things, smartphones, and tablets are generating billions of pieces of data that are instantly uploaded to the cloud. By 2020, the amount of data produced will exceed 44 zettabytes or 44 trillion gigabytes - estimated to be 58 times the amount of all the grains of sand on all the beaches on earth. Enterprise data centers capture, store, process and share this massive amount of data. Cloud computing and big data are driving the transformation of the data center. "Time is Money" used to be the driving factor to differentiate in business, now this is complemented with "Data is Money". The ability to store and process enormous amounts of data are tremendous competitive advantages.

Traditional architectures in data centers usually have three distinct and stand-alone parts:

- Computing with best-in-class processing cores, caches, and attached DRAM memories
- Data storage with lowest cost per bit SSDs or HDDs
- Networking to interconnect the computing part and data storage part

A growing market trend in data centers, called the hyper-converged server, is where computing, storage and networking are fully integrated and condensed in a compact form-factor. Based on some recent reports, 40% of enterprise data centers already use hyper-converged units and this market is expected to grow close to 80% by 2020. The

hyper-converged server, bringing all components in the same box, can now be interconnected in a more efficient way, delivering high capacity, reducing latencies and total cost of ownership, and saving power.

Data center infrastructure's rapid evolution requires high performance, low latency, high capacity, scalable storage solutions optimized in response to elastic demand. Over the last few years, Intel-based server CPU computing has evolved and runs at fast speeds. However, the storage technology in servers has not evolved at the same pace, suffering from low IOPS and high latency, becoming a system performance bottleneck.

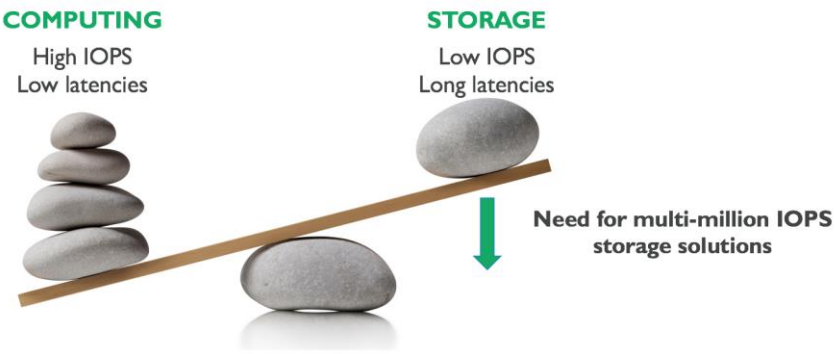

Current performance bottleneck is on storage side

NAND Flash program operation is slow and done at the granularity of a large page size. Current 3D MLC/TLC NAND Flash need about one millisecond to program 16Kbytes page. NAND Flash has to be erased prior to being programmed. The NAND erase operation is slow, in the 10ms range, and is done for a very large block size, 4-8Mbytes.

The Garbage Collection is an additional layer of data management required to properly free-up blocks with obsolete data when storage is idle. The problem occurs when a new request comes while the garbage collection is moving data from blocks to blocks. This typically introduces long and undeterministic latencies in the range of seconds.

Consequently, write operations from host CPU to SSDs generally involve more than one internal write operation between the SSD controller, NAND Flash and DRAM components. This disparity is known as write amplification (WA), most systems typically have WA somewhere between 3 and 4. Obviously, write amplification is undesirable because it means that more data is being written to the media, increasing wear and negatively impacting performance by consuming precious bandwidth to the Flash memory. Flash limitations get harder to hide in SSDs with write amplification and write cliff, garbage collection, and read latency at small queue depth weakness.

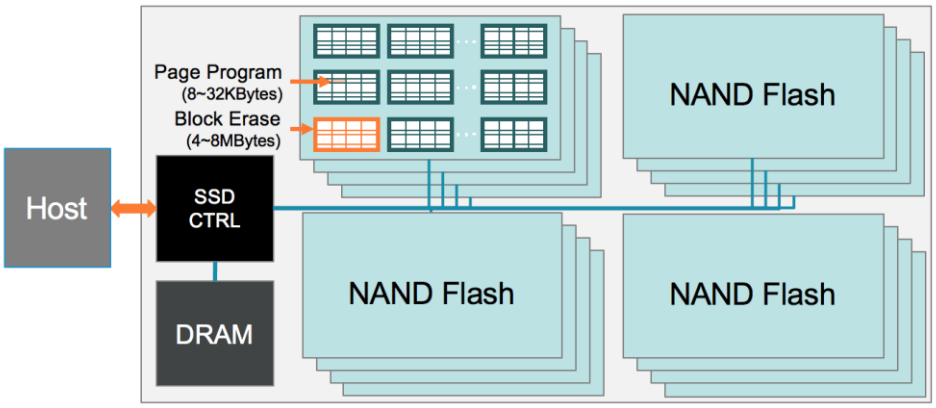

Fresh Out-of-Box performance is not representative of performance under constant load

Compared to today's best-in-class 3D NAND Flash memory, Crossbar's ReRAM memory technologies deliver 100X lower read latency, 1000x faster write performance, 2x higher capacity, 20x lower power consumption.

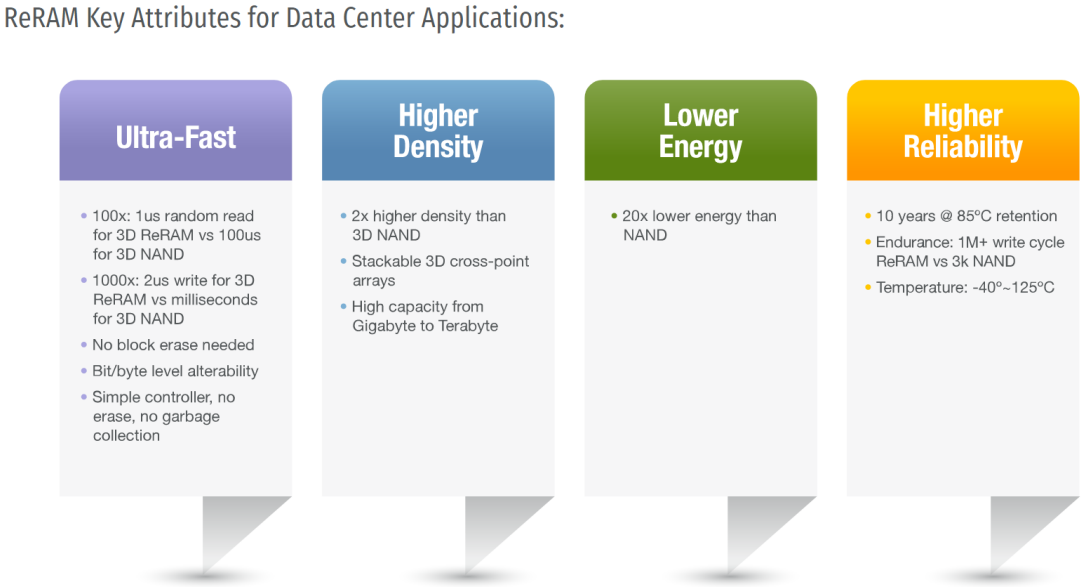

From a storage solution viewpoint, ReRAM technologies don't have erase operation and garbage collection, support bit/byte overwrite capability, can be architected with small pages (e.g. 256Byte pages vs. 16KByte pages in NAND), reduce complexity of the storage controller, and provide impressive performance boosts over Flash-based Solid-State solutions. The ReRAM-based NVMe SSD storage solution read/write latency is 10x lower than Flash-based NVMe SSD; the IOPS @ 512B IO is 8x faster than Flash NVMe SSD; the ReRAM-based NVDIMM read/write latency is even 25x lower than Flash-based NVMe SSD; and the IOPS @ 512B IO is 32x faster than Flash NVMe SSD.

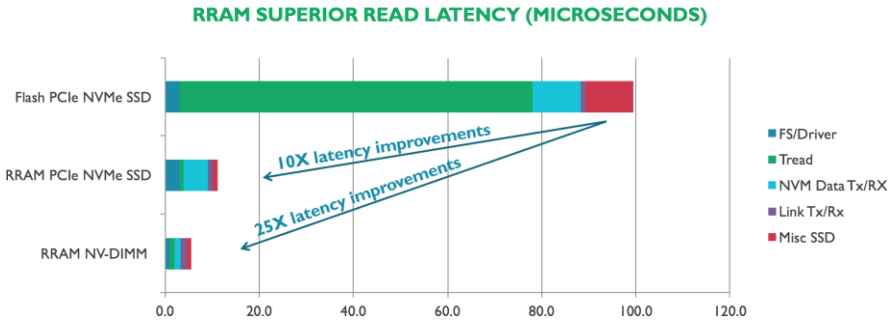

ReRAM solving storage latency limitations

ReRAM technologies deliver even smaller IO size, redefining the way storage solutions are accessed by host CPUs. Correspondingly, the performance gap between storage technologies and computing can be dramatically reduced. Server and data center architectures can be re-invented with ReRAM technology, enabling brand new smaller IO size architectures, paving the way to a new era of tightly integrated memory and computing systems.

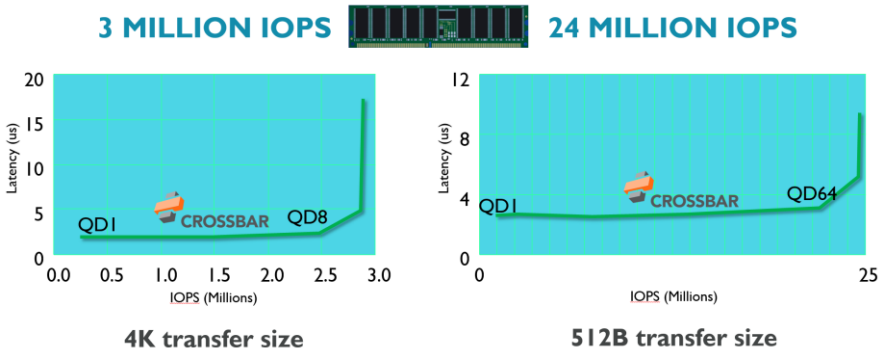

Redefining storage solutions with smaller IO transfer size

To summarize, Crossbar ReRAM technologies revolutionize NVM technologies, answer the density, performance and latency challenges of data centers and storage solutions, significantly reduce the performance gap between storage and computing, enable new system architectures for high performance, lower power, lower costs, and launch a new era of data center storage solutions.