NEWS RELEASE

Contact:

Erin Hartin
HGST
Office: 303-284-7790
Erin.Hartin@HGST.com

Katie Watson
Porter Novelli
Cell: 408-439-2002
HGST_Team@porternovelli.com

HGST RESEARCH DEMONSTRATES BREAKTHROUGH
PERSISTENT MEMORY FABRIC AT FLASH MEMORY SUMMIT 2015

Live Demo Showcases Low Power DRAM Alternative with Greater Scalability for In-memory Compute Applications

Santa Clara, CA, Flash Memory Summit 2015 (HGST Booth #645-647) - August 10, 2015 - Building upon last-year's record-breaking three million I/O per second Phase Change Memory (PCM) demonstration[1], HGST, a Western Digital Company, in collaboration with Mellanox Technologies, is showcasing a revolutionary PCM-based, RDMA-enabled in-memory compute cluster architecture that delivers DRAM-like performance at a lower cost of ownership[2] with greater scalability.

In-memory computing is one of today's hottest data center trends. Gartner Group projects that software revenue alone for this market will exceed US $9B by the end of 2018[3]. In-memory computing enables organizations to gain business value from real-time insights by offering faster performance and greater scalability than legacy architectures.

While modern data center applications can benefit from more main memory, today's DRAM approaches are expensive to scale because of that memory's volatility: DRAM stores data in leaky capacitors, and thus needs to be rewritten many times per second to stave off data loss. This refresh power consumption can be as much as 20-30% of the total server energy[4]. Emerging non-volatile memory technologies, such as PCM, do not have this refresh power demand thereby enabling far greater scalability of main memory than DRAM.

HGST's breakthrough persistent memory fabric technology delivers reliable, scalable, low-power memory with DRAM-like performance, and does not require BIOS modification nor rewriting of applications. Memory mapping of remote PCM using the Remote Direct Memory Access (RDMA) protocol over networking infrastructures, such as Ethernet or InfiniBand, enables a seamless, wide scale deployment of in-memory computing. This network-based approach allows applications to harness the non-volatile PCM across multiple computers to scale out as needed.

The HGST/Mellanox demonstration achieves random access latency of less than two microseconds for 512 B reads, and throughput exceeding 3.5 GB/s for two KB block sizes using RDMA over InfiniBand.

"DRAM is expensive and consumes significant power, but today's alternatives lack sufficient density and are too slow to be a viable replacement," said Steve Campbell, HGST's chief technology officer. "Last year our Research arm demonstrated Phase Change Memory as a viable DRAM performance alternative at a new price and capacity tier bridging main memory and persistent storage. To scale out this level of performance across the data center requires further innovation. Our work with Mellanox proves that non-volatile main memory can be mapped across a network with latencies that fit inside the performance envelope of in-memory compute applications."

---

[1] https://www.hgst.com/press-room/press-releases/HGST-Research-Demonstrates-World-s-Fastest-SSD-at-Flash-Memory-Summit-
[2] http://thememoryguy.com/how-nand-flash-can-reduce-dram-requirements/#more-1138
[3] Market Guide for In-Memory Computing Technologies. Gartner Group, September, 2014.
[4] http://www.itrs.net/ITRS%201999-2014%20Mtgs,%20Presentations%20&%20Links/2013ITRS/Summary2013.htm

"Mellanox is excited to be working with HGST to drive persistent memory fabrics," said Kevin Deierling, vice president of marketing at Mellanox Technologies. "To truly shake up the economics of the in-memory compute ecosystem will require a combination of networking and storage working together transparently to minimize latency and maximize scalability. With this demonstration, we were able to leverage RDMA over InfiniBand to achieve record-breaking round-trip latencies under two microseconds. In the future, our goal is to support PCM access using both InfiniBand and RDMA over Converged Ethernet (RoCE) to increase the scalability and lower the cost of in-memory applications."

"Taking full advantage of the extremely low latency of PCM across a network has been a grand challenge, seemingly requiring entirely new processor and network architectures and rewriting of the application software," said Dr. Zvonimir Bandic, manager of Storage Architecture at HGST Research. "Our big breakthrough came when we applied the PCI Express Peer-to-Peer technology, inspired by supercomputers using general purpose GPUs, to create this low latency storage fabric using commodity server hardware. This demonstration is another key step enabling seamless adoption of emerging non-volatile memories into the data center."

The Persistent Memory Fabric will be demonstrated in the HGST booth #645-647 at the 2015 Flash Memory Summit (http://www.flashmemorysummit.com/) in the Santa Clara Convention Center, Santa Clara, CA from August 11-13, 2015.

About HGST
HGST, a Western Digital company (NASDAQ: WDC), develops innovative, advanced hard disk drives, enterprise-class solid state drives, external storage solutions and services used to store, preserve and manage the world's most valued data. HGST addresses customers' rapidly changing storage needs by delivering intelligent storage devices that tightly integrate hardware and software to maximize solution performance. Founded by the pioneers of hard drives, HGST provides high-value storage for a broad range of market segments, including Enterprise, Cloud, Datacenter, Mobile Computing, Consumer Electronics and Personal Storage. HGST was established in 2003 and maintains its U.S. headquarters in San Jose, California. For more information, please visit the company's website at http://www.hgst.com.