ABSTRACT EXAMPLES

1. NVMe based storage networks can be expanded to thousands of drives with high speed connections such as 40 GbE. A new approach eliminates the need for controller-based storage arrays and is actually faster than internally installed NVMe. The result is very high-performance, networked storage which can take advantage of both flash and Intel Optane storage class memory. Specific case studies demonstrate the real-world TCO/TOI advantages that such networked NVMe-based systems bring to enterprises.
2. Persistent memory, a new technology which offers non-volatility on the memory bus, is emerging as a major approach to improving application performance. Several technologies are under development, including 3D Xpoint, STT-MRAM, PCM, and ReRAM. They allow a wide range of capacity, performance, and price points for computer systems. A method to evaluate and test persistent memory solutions running at-speed is essential to fully understand their benefits and limitations. A low-latency flexible prototyping platform (called Con Tutto) is currently in use for IBM Power servers to evaluate technology characteristics and application benefits. Sample results show the advantages of persistent memory for a MongoDB application using NVDIMM-N.
3. NVMe over Fabrics (NVMe-oF) allows designers to network large numbers of NVMe-based storage arrays. Though this approach addresses capacity scaling, the resulting data transfers are CPU-intensive and hence can bog down the entire solution. A promising solution is to use hardware accelerators to offload the protocols for both NVMe-over-fabrics and RDMA. An even better approach goes beyond the protocols and adds custom and reconfigurable accelerators in the data path. This approach eliminates the CPU bottleneck entirely. Example accelerators utilize FPGAs to achieve significant acceleration while reducing costs.
4. 3D NAND flash memory has been adopted widely in storage systems to lower cost per bit and improve system performance. However, the physical properties and operating methods of 3D NAND differ from predecessor technologies due to differences in materials and structure. Consequently, many reliability issues occur, and novel controller methodologies are required to overcome them. Machine learning can be used to improve the performance and reliability of storage systems based on 3D NAND flash memory. Experimental results based on a machine learning algorithm show read performance improved by 30x, and life cycle enhanced by 3x as a result of optimizing the data structure, error handling methodology, and NAND flash operating methods.
5. New technological developments in NAND flash memory, such as 3D flash and MLC flash, have led to new tradeoffs with regard to such characteristics as endurance and reliability. Currently controller software must be rewritten to account for the changes. An adaptive, self-learning algorithm could avoid the need for such revisions. It would adapt to internal or external parameters that define host behavior, system data structures, and NAND flash condition. It would feed the parameter values into the algorithm and thus determine the best approach. The advantages include simpler FTL software; optimized, data-driven decisions made automatically as the system learns along the iterative process; maximum system performance and reliability; and no need to upgrade software to account for new technologies. Note that manufacturers would also no longer have to consider multiple software versions when performing support, maintenance, and testing.

BIOGRAPHY EXAMPLES

- Manuel d'Abreu is co-founder and Sr VP R&D at Smart IOPS. He was previously a senior Fellow at SanDisk, where he worked on memory management and error correction. Dr. d'Abreu has over 35 years of experience in the semiconductor industry. His expertise spans system-on-chip and system-in-package design, test and manufacturing disciplines, and signal processing, as well as flash storage technology and coding theory. He is also a special advisor to the National Science Foundation. His work has been published at more than 50 technical conferences and in industry journals. He has been granted 45 patents and is credited with 18 pending ones. He earned a PhD from the University of Texas at Austin.

- Andy Mills is the CEO and co-founder of Enmotus, a storage automation and analytics software company. He heads the development of a high performance next generation SSD and storage device level virtualization and automigration technology. He was previously VP Marketing and Business Development at Dot Hill Systems, a storage systems manufacturer. He was in charge of all of Dot Hill's products and led the development of a PERC RAID controller solution for Dell. He has also held executive management positions at Ciprico, TDK, Rockwell, and AMD. He holds 8 patents related to storage and networking, has published a book with Prentice-Hall, and has given presentations at many conferences including Flash Memory Summit. He earned a Master's of Engineering in electrical and computer engineering and a Bachelor's of Engineering (Honors) from Bangor University (Wales, UK).

- HK Verma is a Principal Engineer at Xilinx developing FPGA-based data center accelerator solutions. He has been working with leading hyperscale customers to integrate optimized FPGA acceleration into their database and storage solutions. He has pioneered successful accelerator adoption on existing RDBMS and NoSQL database storage solutions by working closely with partners to define their future work. He has also been co-founder and VP at Velogix, a startup offering programmable compute silicon and software. He holds 34 issued US patents and has presented tutorials and papers at leading conferences. He holds an MSEE from UCSB and a BTech from IIT Madras in electrical engineering.

- Adrian Jackson is a Research Architect at EPCC, a leading European supercomputing center located at the University of Edinburgh. He currently leads the Architecture work package for the European funded NEXTGenIO project, which is designing an innovative hardware and software solution for I/O at the peta-and exascale system level using NVRAM. He has presented at several HPC-related conferences and has published articles in IEEE Transactions on Parallel & Distributed Systems and the Journal of Computer Science Research and Development. He earned an MSc in high-performance computing from EPCC and a bachelor's degree in computer science from the University of Edinburgh.

- Rob Davis is VP Storage Technology at Mellanox where he focuses on applying their high-speed interfaces (such as 40G and 100G) to storage systems. Over the last few years he has moved Mellanox into a leadership position in NVMe over Fabrics (NVMe-oF). As a technology leader and visionary for over 35 years, he has been a key figure in the development of an entire generation of storage networking products. Davis was previously VP/CTO at QLogic, where he drove development and marketing of Fibre Channel, Ethernet, and InfiniBand technology into new areas such as blade servers. Before joining QLogic, Davis worked at Ancor Communications, where he drove development and marketing of Fibre Channel and InfiniBand products.

TITLE EXAMPLES

- Predicting SSD Performance for Today's Dynamic Workloads
- Scaling NVMe Storage Networks to Thousands of Drives
- Using Machine Learning to Enhance Flash Endurance and Throughput
- Using FPGAs to Accelerate NVMe-oF-Based Storage Networks
- A Self-Learning Flash Translation Layer (FTL) for NAND Flash Controllers