



Excelero and ThinkParQ Go to Market with Integrated NVMe[®] and BeeGFS

In benchmark testing at early deployments, combined solutions deliver 3X performance vs. direct-attached NVMe, helps teams maximize data center efficiency

SAN JOSE CA and KAISERSLAUTERN, GERMANY May 28, 2019 – [Excelero](#), a disruptor in software-defined block storage, and [ThinkParQ](#) announced that their joint go-to-market relationship is bearing fruit and the firms have successfully integrated Excelero's NVMe[®] software-defined block storage solution for shared NVMe at local latency along with ThinkParQ's leading parallel cluster file system BeeGFS. Benchmark results from early deployments spotlight the maximum resource utilization and flexibility that are possible from combining the two solutions – performance that is perfectly suited for data-intensive workloads such as high performance computing (HPC), AI, ML, and analytics.

In response to customer demand, Excelero and ThinkParQ already have deployed NVMe storage clusters with BeeGFS and NVMe to increase productivity of advanced GPU compute clusters. These organizations experienced the storage bottleneck issue endemic to AI – that I/O intensive applications cannot be processed fast enough even on GPU-based systems. As a result, IT teams often copy datasets in and out locally, which is an inefficient and time-consuming process.

“With Excelero's NVMe, our customers have access to an ultra-low latency, high performance approach to scale-out storage” said Frank Herold, CEO of ThinkParQ. “We've been impressed with NVMe's ability to deliver the high IOPS and ultra-low latency of NVMe drives over the network with highly available volumes – as well as options for distributed erasure coding and BeeGFS' unmatched ability to efficiently handle all kinds of access patterns and file sizes.”

To demonstrate the possibilities of this new scale-out infrastructure, the companies used the industry standard mdtest and IOR benchmarks. The test system was a compact 2U 4-server chassis with a total of 24 NVMe drives, connected via a 100Gbit RDMA network to 8 BeeGFS client compute nodes. Tests were run on the exact same hardware configuration with BeeGFS utilizing the direct-attached NVMe vs. BeeGFS utilizing NVMe logical volumes. Taking advantage of NVMe to offload mirroring operations, BeeGFS file create operations were boosted 3x, while metadata read operations were boosted 2.5x. For small random file access, which is often considered to be especially critical for application efficiency, NVMe's low latency technology boosted BeeGFS 4K write IOPS to 1.25 million per second, a 2.5x improvement. By leveraging NVMe distributed erasure coding for BeeGFS, customers can get up to 90% usable capacity while still tolerating drive failures. All the while, achieving 75GB/s streaming throughput from this entry-level system.

Through June 30 Excelero is offering enhanced pricing on combined purchases of NVMe and BeeGFS. More details are available from both companies.

“You can't build tomorrow's enterprise on yesterday's infrastructure,” said Lior Gal, CEO and co-founder of Excelero. “We're delighted at how smoothly NVMe and BeeGFS work together, and the way we enable organizations to work without compromise to either IT teams or end users. We look forward to working further with ThinkParQ so that more organizations worldwide can maximize their NVMe ROI and the efficiency of their entire operations.”

Excelero and ThinkParQ will showcase the performance and efficiency of BeeGFS powered by NVMe at the ISC High Performance Event 2019, June 16-20 in Frankfurt in ThinkParQ's booth J-640 and Excelero's booth E-1039. For more information, visit www.thinkparq.com or www.excelero.com

About ThinkParQ GmbH

[ThinkParQ](#) strives to create and develop the fastest, most flexible, and most stable solutions for every performance-oriented environment. Established in 2014 as a spinoff from the Fraunhofer Center for High-Performance Computing, ThinkParQ drives the research and development of BeeGFS, and works closely with system integrators to create turn-key solutions.

Visit <http://www.thinkparq.com> for further information.

About Excelero

Excelero delivers low-latency distributed block storage for web-scale applications such as AI, machine learning and GPU computing. Founded in 2014 by a team of storage veterans and inspired by the Tech Giants' shared-nothing architectures for web-scale applications, the company has designed a software-defined block storage solution that meets the low-latency performance and scalability requirements of the largest web-scale and enterprise applications.

Excelero's NVMeMesh enables shared NVMe across any network and supports any local or distributed file system. Customers benefit from the performance of local flash with the convenience of centralized storage while avoiding proprietary hardware lock-in and reducing the overall storage TCO. NVMeMesh is deployed by major web-scale customers, for data analytics and machine learning applications and in Media & Entertainment post-production and HPC environments.

Follow us on Twitter @Excelerostorage, on [LinkedIn](#) or visit us at www.excelero.com to learn more.