

44GB/s throughput on the Boston Flash-IO Talyn – NVMeOF Solution

The latest impressive benchmarks to come out of the renowned Boston Labs on the Boston Flash-IO Talyn solution powered by Excelero NVMesh.

FRANKFURT, MUNCHEN, GERMANY, June 25, 2018 /EINPresswire.com/ -- Frankfurt Germany, – 25th June, 2018 - [Boston](#) are pleased to announce the latest impressive benchmarks to come out of the renowned Boston Labs on the Boston [Flash-IO Talyn](#) solution powered by [Excelero](#) NVMesh. The solution delivered a sequential read of 44GB/s using 4MB blocks and a random read of 840k IOPS using 4k blocks.

The solution that was used consisted of four servers with each server including 4 NVMe drives – 16 in total.

Patrick O’Neill, HPC Systems Engineer, Boston Labs says “Using an NVIDIA® DGX-1™ as a client during testing we were able to achieve a throughput of 44GB/s for sequential reads and 840k IOPs for random reads. Both figures demonstrate the impressive performance that can be realised when using Flash-IO Talyn under varying workloads, even when only using a single client setup.”

“

Using an NVIDIA DGX-1 as a client during testing we were able to achieve a throughput of 44GB/s for sequential reads and 840k IOPs for random reads.”

Patrick O’Neill, HPC Systems Engineer, Boston

solutions.”

The Boston Flash-IO Talyn will be demonstrated with Excelero alongside an NVIDIA® DGX-1® at ISC 2018 this week on stand [C-1232].

This solution, based on Supermicro, and powered by Micron and Mellanox, represents a revolutionary



Excelero Logo



Flash IO Talyn

Yaniv Romem, CTO for Excelero adds; “Modern GPUs used in AI and ML have an amazing appetite for data - up to 16GB/s per GPU. Starving that appetite with slow storage, or wasting time copying data back and forth wastes the most precious (expensive) resource you've purchased. Talyn is amazing because it gives you a building block to feed virtually any size NVIDIA GPU farm with scalable simplicity. The combination of NVMesh performance on the optimized Talyn hardware platform gives you the affordability of commodity hardware but with the ease of deployment of proprietary

leap forward for NVMe over fabrics extending the promise of SDS to low latency workloads by leveraging server-side NVMe-based flash storage to deliver a scalable converged infrastructure for next level performance. In a world where data is growing rapidly and the need to capture, present and transform must be reconsidered, the main aim of the system is to accelerate the data feed to your

NVIDIA GPU servers ensuring your GPU can operate to its full potential.

Designed as a dynamic block to offer a cost-effective way of testing features and performance of NVMe over fabrics technically applying Excelero NVMesh, the timing could not be more perfect in an era where there's major focus on deep learning solutions and how they are processing data faster than typical storage bandwidth, thus a fitting feature for a show focused on the drive of HPC solutions.

The NVIDIA DGX-1 is an AI supercomputing system powered by eight of the world's most advanced datacentre GPUs – the NVIDIA® Tesla® V100 with Tensor Core architecture, and, incorporating next-generation NVIDIA NVLink™. DGX-1 is purpose-built for the unique demands of AI and deep learning and leverages the NVIDIA GPU Cloud Deep Learning Software Stack to deliver maximized GPU accelerated performance.

ENDS

For further information please contact:

Maz Lopez, Head of Marketing

Boston Limited

+44 (0) 1727 876 100

maz.lopez@boston.co.uk

Tom Leyden

Vice President Corporate Marketing

+1 818 919 84 78

+32 473922211

tom.leyden@excelero.com

Skype: tomleyden

Notes to editors:

1. About Boston

Boston are the highest accredited NVIDIA Elite Solution Providers in the UK and are delivery partners of the NVIDIA Deep Learning Institute.

Boston Limited has been providing cutting-edge technology since 1992 using Supermicro® building blocks. Our high performance, mission-critical server and storage solutions can be tailored for each specific client, helping you to create your ideal solution. From the initial specification, solution design and even full custom branding – we can help you solve your toughest business challenges simply and effectively. Follow Boston on Twitter @BostonLimited, on LinkedIn or at Boston.co.uk to learn more.

2. About Excelero

Excelero enables enterprises and service providers to design scale-out storage infrastructures leveraging standard servers and high-performance flash storage. Founded in 2014 by a team of storage veterans and inspired by the tech giants' shared-nothing architectures for web-scale



Servers | Storage | Solutions

Boston Logo

applications, the company has designed a software-defined block storage solution that meets performance and scalability requirements of the largest web-scale and enterprise applications.

With Excelero's NVMesh, customers can build distributed, high-performance server SAN for mixed application workloads. 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. The solution has been deployed for hyper-scale Industrial IoT services, machine learning applications and massive-scale simulation visualization.

Follow Excelero on Twitter @Excelerostorage, on LinkedIn or visit us at www.excelero.com to learn more.

Mary Lopez
Boston Limited
+44 (0) 1727316129
email us here

This press release can be viewed online at: <http://www.einpresswire.com>

Disclaimer: If you have any questions regarding information in this press release please contact the company listed in the press release. Please do not contact EIN Presswire. We will be unable to assist you with your inquiry. EIN Presswire disclaims any content contained in these releases.

© 1995-2018 IPD Group, Inc. All Right Reserved.