



ISC High Performance
The HPC Event.

ROUND UP

FOREWORD


International Supercomputing Conference, or ISC High Performance as it's better known, is in its 33rd year as the premiere European event for HPC. Once again held in Frankfurt, the conference was a record-breaking one, attracting over 3,400 attendees and 162 exhibitors, including Boston.

Held for the third consecutive year at the Messe Frankfurt, the event ran between 24th-27th June with the main action taking place on the exhibition floor.

Unsurprisingly, the theme of the event was focused around high performance computing, networking and storage; with exhibitors and attendees travelling from around the world to attend.



There was good news for consumers, with more choice than ever of CPU architecture with POWER9 and ARM featured in the mix, but of course, there was no getting away from the fact that a lot of attention was focused on the Green Team with plenty of exciting demonstrations and announcements from NVIDIA and their partners.



In this newsletter we explore some of the key-technologies that caused a stir including; the European debut of the much awaited Supermicro® SuperServer featuring 16x NVIDIA® V100 GPUs on the new NVIDIA HGX-2 platform, as well as the Boston and Asperitas Immersed Computing® collaboration featuring AMD® EPYC™.

MORE ON HGX-2



Announced at Computex, Taipei at the end of May, the world's most powerful cloud server platform for AI and HPC, NVIDIA HGX-2, was on display at ISC; the first time it has made an appearance in Europe. During ISC, there was a steady stream of attendees keen to see the SuperServer in real-life & photograph it. Similarly, to the NVIDIA DGX-2, this SuperServer also features 16 of the latest generation NVIDIA V100 GPUs.

Charles Liang, president and CEO of Supermicro commented at the launch; "To help address the rapidly expanding size of AI models that sometimes require weeks to train, Supermicro is developing cloud servers based on the HGX-2 platform that will deliver more than double the performance."

[MORE ABOUT HGX-2](#)

A NEW #1 SUPER COMPUTER

With almost 25,000 entries to the TOP500 list in the last 25 years, updated twice annually, the TOP500 highlights the 500 most powerful commercially available computer systems known and ranked by their performance on the LINPACK Benchmark.

The June list announcement at ISC delivered a dramatic shake-up at the top, with a new number one system, that meant that the US claimed the most powerful super-computer in the world for the first time since November 2012.



"Summit," running at the U.S. Department of Energy's (DOE) Oak Ridge National Laboratory (ORNL), broke the 100 PF barrier delivering an astounding 122.3 PF on High Performance Linpack (HPL). Summit consists of 4,356 nodes, each equipped with two 22-core IBM POWER9 CPUs and six NVIDIA Tesla V100 GPUs, interconnected with Mellanox dual-rail EDR InfiniBand.

[VIEW TOP500](#)

BOSTON GOES GREEN IN MORE WAYS THAN ONE

At a time where both power and space come at a premium the HPC community is always looking for innovative solutions .

Boston certainly led the way when it came to demonstrations of leading-edge technologies that were the talk of the show.

Some likened it to a cryogenic chamber and others a space ship – Immersed Computing® pioneers Asperitas joined Boston to demonstrate their datacentre-in-a-box solution populated with custom configured sleds just launched from Boston Labs featuring AMD® EPYC™.

Glowing blue with a bright red panic button on the front, there's no denying this solution has style...but it's what's under the hood that's truly impressive...




Compared to an average air-cooled datacentre, Immersed Computing gives 5-10 times traditional hardware density and

can save more than 50% of the total energy footprint. By using immersion, 10-45% of IT energy is reduced due to the lack of fans, while other energy consumers like cooling installations can achieve up to 95% energy reduction.

Did you "dare to dip?" Tweet your picture using #BostonHPC - Brave ISC attendees who "dared to dip" their fingers in the datacentre-in-a-box walked away with a "Don't Panic" towel.

As NVIDIA's largest and most-accredited Elite Partner in Northern Europe, and as global delivery partners of the NVIDIA Deep Learning Institute; the Boston stand featured more GPU and Machine Learning solution demonstrations than ever before.

NVMe also began to turn up in a range of solutions, we saw it populated in compute nodes and used as burst buffers, becoming the norm for the top tier of storage in HPC.



Boston Labs announced the latest impressive benchmarks on the Boston Flash-IO Talyn solution, powered by Exceero NVMe. The solution delivered sequential read of 44GB/s using 4MB blocks and random read of 840k IOPS using 4k blocks.

[READ MORE](#)

NVME OVER FABRICS AND DEEP LEARNING ON THE DGX-1

The Flash-IO Talyn was demonstrated at ISC on the Boston stand by Boston Labs and Excelero. 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."



Yaniv Romem, CTO for Excelero added; "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 NVMe performance on the optimized Talyn hardware platform gives you the affordability of commodity hardware but with the ease of deployment of proprietary solutions."

[READ MORE](#)

REAL-TIME HYPERSCALE ERASURE CODING

Nyriad also joined Boston to live demo the Boston Igloo Nebari the world's first "GPU-powered SSD storage solution on Supermicro". The solution provides speed and data resilience exceeding traditional RAID-based solutions.

Supporting up to 72x Micron® 5100 PRO SSDs and Micron NVDIMMs powered by NVIDIA® Tesla® P4 GPUs, the Nebari harnesses Nyriad's NSULATE™; a GPU-accelerated, software-defined alternative to RAID. NSULATE presents groups of physical storage devices as a Linux block device to the system and leverages the accelerated processing capabilities of NVIDIA GPUs to perform real-time erasure coding at many times the resilience and speed of CPU erasure coding or RAID-based systems.

The live demonstration of the Igloo Nebari showed off the high-parity support by allowing over 50% of the SSDs in the array to be removed with no dropped frames when playing a high-resolution video.



ULTRA

Supermicro once again showcased their "Ruler" form factor Ultra server featuring Intel NVMe SSDs offering extreme density and performance storage.

Alongside the "Ruler" was the 1U – 18x 7mm Hot-Swap NVMe Drive Bays.

High-density Ultra servers, storage bridge bay dual ported NVMe servers and JBOF's [Just a Bunch Of Flash] are available in Supermicro's range of NVMe platforms, there is something for every application, regardless of which camp you are in. It can be difficult to navigate the options available and know which the right choice for you is, fortunately Boston Labs are on hand with their expertise to steer you towards the right choice for your needs.

BOSTON
Servers | Storage | Solutions

BOSTON LIMITED,
UNIT 5 CURO PARK, FROGMORE, ST ALBANS, AL2 2DD
TEL: +44(0)1727876100 / WEB: www.boston.co.uk