



HPC BROCHURE

ISSUE 2019, Q2



PROUD TO PARTNER WITH:



Technology
Provider
Platinum 2019



ELITE
PARTNER

CONTENTS



3

SERVICES

3 Services

4

BOSTON LABS

4 HPC Cluster Architecture
5 Boston Labs Introduction
6 Tailor Made & Customised
7 Planning, Installation, Testing, QA
8 Additional Managed Services & Support

9

GPU SOLUTIONS

9 GPU Solutions
10 NVIDIA® DGX™ Family
11 Supermicro® NVIDIA® Family
12 Supermicro® AMD Family

13

STORAGE SOLUTIONS

13 Storage Solutions
14 vScaler® Solution
16 Nyriad Solution
18 BeeGFS® Solution
20 Object Storage Solution

22 Quobyte Solution
24 NVMe/Burst Buffer Solution
26 Bright Computing
28 Immersed Computing®

31

NETWORKING

31 Networking
32 Mellanox Quantum™ HDR Switch
33 Mellanox Spectrum™ Switch

34

HPC NODES

34 Compute Nodes
34 GPU Nodes
34 Storage Nodes
34 Storage Nodes
35 AMD EPYC™
35 GPU Nodes
35 Storage Nodes

36

TECH TALK

36 Know Your Limits...
39 InstaDeep™ Powers AI...
41 Why Won't You Change?...

SERVICES



SERVICES



HPC CONSULTANCY SERVICES

Our HPC experts, based around the world, are constantly benchmarking new solutions, the results of which we share with our customers, to deliver fully optimised solutions.



BOSTON LABS

Boston has, in excess of, \$2m of hardware that's available to customers, either onsite or via remote access, to test out the latest technologies and architectures. Boston customers can run their own workloads, benchmarks and simulations on configured systems prior to purchase.



TAILOR MADE SOLUTIONS

Boston deliver bespoke solutions matched to customers' applications and requirements. Through the utilisation of hardware and software accelerators, our solutions are designed around upgraded performance and faster workflows.



PLANNING & INSTALLATION

Our senior engineers understand the technical dependencies and requirements of your organisation, we will ensure a well thought out installation is managed and completed on schedule and with the utmost professionalism.



QUALITY ASSURANCE

Boston has numerous state-of-the-art build facilities, that are managed by our in-house engineering teams, to ensure all our solutions are built, and configured, to the highest standards for quality, stability and performance.



ADDITIONAL MANAGED SERVICES & SUPPORT

From tailor-made SLA solutions, warranty support and spares packages – Boston offers customised services level work packages for the ongoing support and management of your solutions.

HPC CLUSTER ARCHITECTURE AND WORKFLOW

When it comes to the High Performance Computing (HPC) market, clusters are rapidly reshaping it, driven in the most part by price/performance and the open source Linux OS.

Even though there are endless opportunities that clusters offer, and clusters, comprised of commodity server hardware and software are gaining acceptance...getting a cluster running, and learning how to use it, requires skilled resources, services and time.

If not installed correctly, this can lead to badly planned clusters with software that hasn't been chosen appropriately.

The result? IT departments turn to costly SMP alternatives because of the standardisation shortcomings of cluster computing.

Boston understand the end-to-end process of building, implementing and managing clusters, so cost is minimised, as is the time to get a cluster fully operational.

Boston offers a range of both open source and commercial cluster management packages, that are selected in-line with our clients environment and workflow.

We are able to pre-engineer and test multiple variations of hardware and software packages in our Boston Lab facilities - making our Linux compute clusters easy to deploy, simple to use, consistent, transparent, turnkey and available.

01 **INGEST**

02 **PARALLELISE**

03 **CLUSTER
MANAGEMENT**

04 **PRIVATE CLOUD**

05 **SCHEDULERS**

06 **SECURITY**

07 **CONTAINERISATION**



WORLD-CLASS BOSTON LABS

The first Boston Labs facility was launched in 2005 just outside London, when it became apparent that our customers were not able to gain a full understanding of latest generation hardware and its application against their workloads/use cases.

Following the success of the UK Lab, further facilities were opened in Germany and India, with future sites earmarked for the future.

01

LEVERAGE OUR EXPERTISE

Understandably, some customers are unable to visit our labs in person. Therefore, arrangements can be made for test systems to be set up in any of our labs with remote access via RDP, SSH or other means; enabling trials on Boston solutions from wherever you are in the world.

02

EXPERT HELP EVERY STEP OF THE WAY

03

AT YOUR CONVENIENCE WITH ANYTIME REMOTE DIAL-IN ACCESS

The Boston Labs facilities have grown in popularity to such an extent that even leading hardware manufacturers and industry journalists utilise them for benchmarking and cluster configuration.

04

BE THE 1ST TO ACCESS THE LATEST AND FUTURE TECHNOLOGIES

To find out the latest technologies available for testing please visit our website.

05

ON-SITE TESTING AVAILABLE



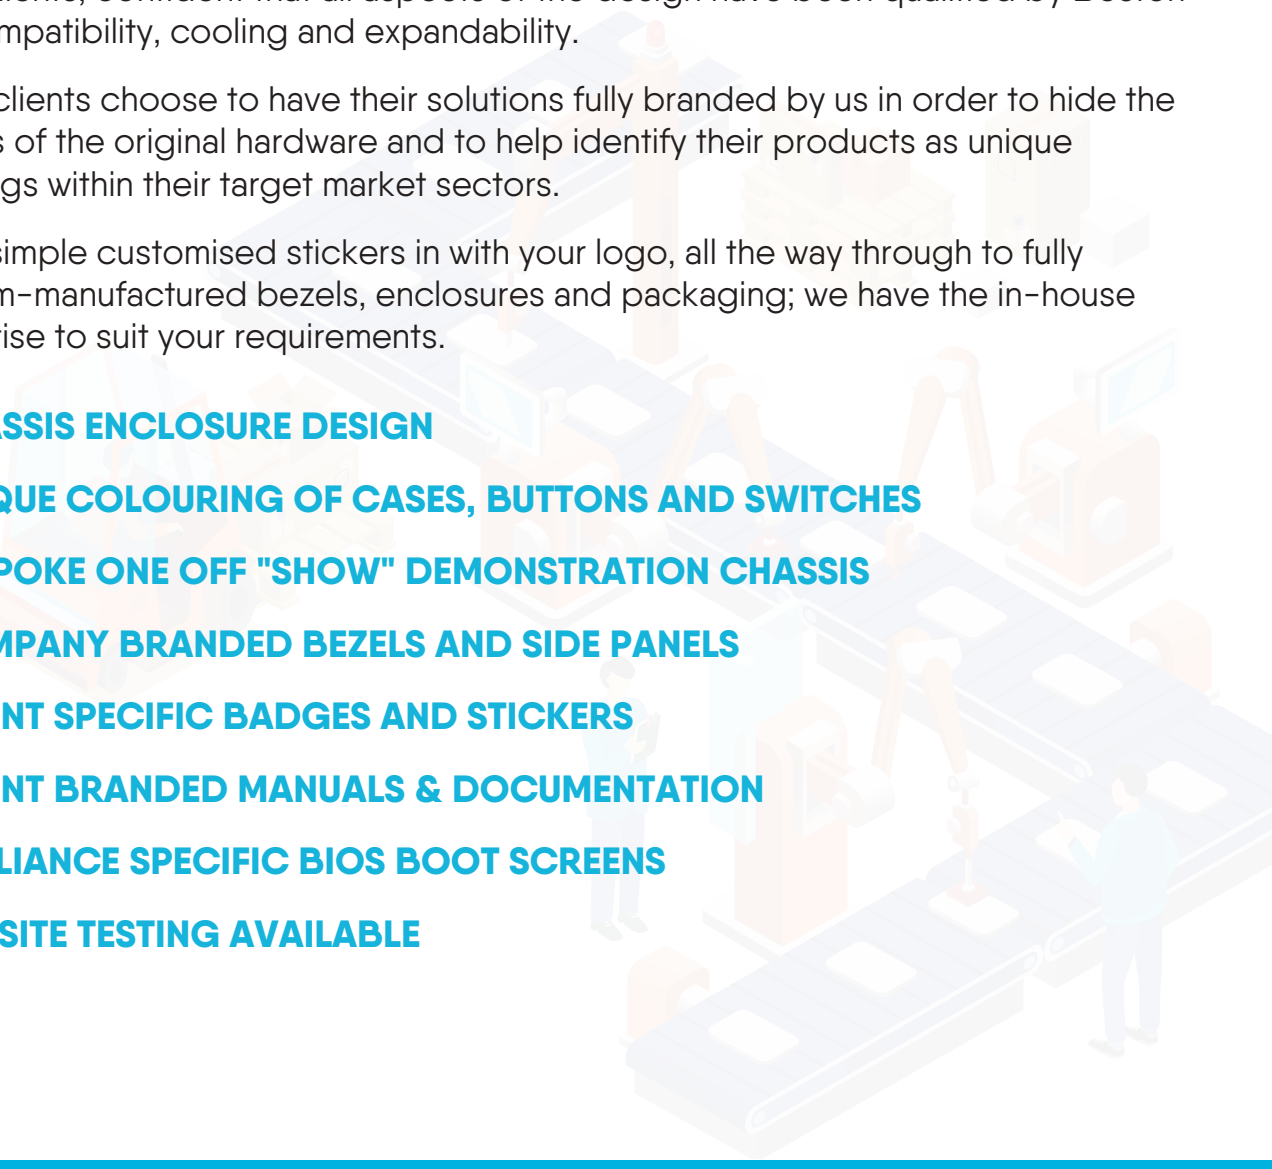
TAILOR MADE & CUSTOMISED

As the chosen OEM manufacturing partner for numerous high-profile brands, we have a great deal of expertise in customising our solutions to meet your exact requirements.

Boston offers custom configuration of a server, workstation or solution needed by your clients, confident that all aspects of the design have been qualified by Boston for compatibility, cooling and expandability.

Many clients choose to have their solutions fully branded by us in order to hide the origins of the original hardware and to help identify their products as unique offerings within their target market sectors.

From simple customised stickers in with your logo, all the way through to fully custom-manufactured bezels, enclosures and packaging; we have the in-house expertise to suit your requirements.

- 
- **CHASSIS ENCLOSURE DESIGN**
 - **UNIQUE COLOURING OF CASES, BUTTONS AND SWITCHES**
 - **BESPOKE ONE OFF "SHOW" DEMONSTRATION CHASSIS**
 - **COMPANY BRANDED BEZELS AND SIDE PANELS**
 - **CLIENT SPECIFIC BADGES AND STICKERS**
 - **CLIENT BRANDED MANUALS & DOCUMENTATION**
 - **APPLIANCE SPECIFIC BIOS BOOT SCREENS**
 - **ON-SITE TESTING AVAILABLE**

PLANNING, INSTALLATION, TESTING & QA

Boston's in-house validation, build, test and QA procedures are second to none; however, we can augment these to reflect your specific needs; including full rack-scale design.

Depending on your circumstances, our expert solution architects can be available to work with you right from the early design phase, all the way through to the delivery and installation of your solution.

01 CONSULT

02 DESIGN

03 ASSEMBLE

04 CONFIGURE

05 TEST

06 SHIP

07 INSTALL

08 SUPPORT

With over 25 years in the industry, and access to leading-edge technologies such as immersed-computing; and leading software vendors, Boston are uniquely placed to offer a full rack-scale design that includes BIOS, firmware, IPMI, OS and customer imaging – including ongoing cluster management.

For more information on Boston's rack-scale design service visit bit.ly/BostonServices



ADDITIONAL MANAGED SERVICES & SUPPORT

Once you have received your new system, our after sales support team are on hand to answer any queries or problems that you may have.

Every support engineer is trained to deal with requests quickly and effectively, using escalation procedures where necessary to ensure maximum up-time. In addition to our standard warranty, we can offer fully tailored service solutions should this be required.

HPC CONFIGURATION CASE STUDY

Challenges:

The UK's Science and Technology Facilities Council (STFC) needed to upgrade the storage infrastructure for the JASMIN super data cluster. With demand for growth of up to 300PB in the next few years.

Solution:

BostonLabs liaised with STFC to understand their specific requirements before configuring a solution that utilised Quobyte's Data Centre File System.

Advantage:

In addition to having S3 connectivity, the Quobyte system also afforded the JASMIN administrators the ability to scale their storage capacity (currently at 42PB) & performance linearly while providing ease of management.

01

**24/7 PHONE
SUPPORT**

02

**ONSITE WARRANTY
SUPPORT**

03

**NEXT DAY AND
SAME DAY
BUSINESS SERVICES**

04

**GLOBAL WARRANTY
COVERAGE
AVAILABLE**

GPU SOLUTIONS



GPU SOLUTIONS FOR HPC



NVIDIA® DGX™ SYSTEMS

As the most accredited NVIDIA Elite Partner, Boston are pleased to offer NVIDIA DGX systems available for purchase, lease and testing along with training from the Deep Learning Institute.



SUPERMICRO® WITH NVIDIA® TESLA® V100

As NVIDIA Tesla Partner of the Year, Boston are pleased to offer a comprehensive catalogue of products that are fully compatible with the full feature-set of NVIDIA Volta GPUs.



SUPERMICRO® WITH AMD RADEON™

As AMD Elite Partners, Boston are experienced in configuring that accommodate both AMD Radeon Instinct™ and AMD Radeon™ Pro GPUs.

GPU SOLUTIONS

NVIDIA® DGX™ SYSTEMS




NVIDIA® DGX™ Systems are built on the new, revolutionary NVIDIA® Volta™ GPU platform. Combined with innovative GPU-optimised software and simplified management, these fully integrated solutions deliver ground-breaking performance and results.

NVIDIA® DGX™ Systems are designed to give data scientists the most powerful tools for AI exploration—tools that allow you to go from your desk to the datacentre to the cloud from the second you switch your system on.



KEY FEATURES

The NVIDIA® Tesla® V100 Tensor Core is the most advanced datacentre GPU ever built to accelerate AI, HPC, and graphics. It's powered by NVIDIA Volta architecture, comes in 16 and 32GB configurations, and offers the performance of up to 100 CPUs in a single GPU.



DGX-1™

The essential instrument for AI research, designed to accelerate your datacentre and streamline your deep learning workflow. Experiment faster, train larger models, and get insights starting on day one.

DGX STATION™

Built on the same software stack that powers NVIDIA® DGX-1™, the NVIDIA® DGX Station™ is the only personal supercomputer for leading-edge AI development from your desk to the datacentre.

DGX-2™

Break through the barriers to AI speed and scale with NVIDIA® DGX-2™, the first 2 petaFLOPS system that engages 16 fully interconnected GPUs for 10X the deep learning performance.

GPU SOLUTIONS

SUPERMICRO® WITH NVIDIA® TESLA® V100



Frequently first to market with solutions that utilise the latest generation of GPUs from NVIDIA®, Supermicro have a comprehensive range that enables the progression of AI from start-up sized organisations all the way to large HPC clusters.

Each solution also offers fantastic energy efficiency with many PSUs running at Titanium (96%) level. Boston are pleased to offer these solutions on their own or as part of a fully-coordinated cluster design and installation.



THE CASE FOR GPU POWERED HPC:

HPC datacentres need to support the ever-growing computing demands of scientists and researchers while staying within a tight budget. The old approach of deploying lots of commodity compute nodes substantially increases costs without proportionally increasing datacentres performance.

With over 550 HPC applications accelerated—including all of the top 15—all HPC customers can now get a dramatic throughput boost for their workloads, while also saving money.

ANNA VOLTA

Housing up to 4x NVIDIA® Tesla® V100 SXM cards, but available part-populated, this dense yet compact server utilises NVIDIA® NVLINK™ for to 300 GB/s GPU to GPU communication.

ANNA X16

Designed as a cloud platform for AI and HPC, combining 16 NVIDIA® Tesla® V100 32GB SXM3 GPUs connected via NVLink and NVSwitch to work as a unified 2 PetaFlop accelerator to deliver huge computational power.

SYS-1019GP-TT

A single socket server for your HPC cluster, with Intel® Xeon® Scalable CPUs, 6 Hot-swap 2.5" SAS/SATA drive bays along with support for up to two NVIDIA® Tesla® P100/P40/P4/M60 GPUs.

GPU SOLUTIONS

SUPERMICRO® WITH AMD RADEON™



Supermicro® offers AMD Radeon Instinct™ supported SuperServers® that boast excellent price performance for machine intelligence & deep learning. Radeon Instinct™ GPUs are designed to be open from the metal forward and provide higher level of datacentre performance and efficiencies through the ROCm open software platform.

- Based on cutting-edge “VEGA” graphics architecture built to handle big data sets and diverse compute workloads
- 64 nCU compute units to accelerate demanding workloads



KEY FEATURES

- Up to 12.3 TFLOPS of peak FP32 compute performance to speed up compute intensive machine intelligence
- Up to 24.6 TFLOPS of FP16 peak compute performance for deep learning training applications
- State-of-the-art memory technology: 16GB of HBM2 memory with ECC2 and high bandwidth cache controller (HBCC)
- Passively cooled, full-height, dual-slot, 300W TDP board power – designed to fit in most standard server designs
- MxGPU for Virtualised Compute Workloads – drive greater utilisation and capacity in the datacentre
- Advanced Remote Manageability Capabilities, for simplified GPU monitoring in large scale systems

SYS-1029GQ-TNRT

This compact server supports up to 4 Radeon Instinct™ accelerators in just 1U making it an ideal server as part of a larger cluster for HPC, big data analytics and more.

SYS-4029GP-TRT2

Perfect for AI, Big Data analytics, hpc, research lab/ national lab, astrophysics and business intelligence workloads, this versatile system can house up to 10 Radeon Instinct™ MI25 accelerators.

SYS-7049GP-TRT

Housed in a handy workstation form factor, 4x Radeon Instinct™ MI25 accelerators can be fully utilised for desk-top artificial intelligence and deep learning applications.

STORAGE SOLUTIONS



STORAGE SOLUTIONS FOR HPC



HYPER-CONVERGED, VIRTUALISED & CLOUD

We can provision all-flash, hyper-converged appliances at a fraction of the cost of our competitors, as well as a broad range of virtualised storage solutions.



SCALE-OUT

Scale-out parallel file systems and high performance computing storage systems.



NETWORK ATTACHED STORAGE (NAS)

A range of network attached file systems and scale out NAS solutions with all the features your enterprise would need.



BURST BUFFER STORAGE

High bandwidth, low latency, all-flash tier for I/O intensive applications on your cluster.



STORAGE CLUSTER MANAGEMENT

Automate and control your storage infrastructure under a single management interface.



CUSTOM SOLUTIONS

Have a specific storage solution in mind? Get in touch and we will tailor make your solution to fit your exact requirements.

CLOUD SOLUTION

MULTI-VERTICAL PRIVATE CLOUD



With vScaler, you can deploy your on-premise private cloud in minutes, add cloud-based services and applications, scale to public and hybrid cloud environments on-demand, all under a single management portal. vScaler simplifies datacentre infrastructure by integrating HPC, Big Data and Cloud resources into a converged platform that is

capable of running any application at any scale. The platform enables agile sysadmin teams to quickly deploy scalable, production-ready private cloud environments, big data analytical platforms and/or traditional HPC clusters all under one management portal.

 **SCALER**

 **SUPERMICR**

KEY FEATURES

Finance, Government, Broadcast & Media, Scientific Research, Oil & Gas, Manufacturing and Bioinformatics are simply a few of the verticals that can leverage the vScaler cloud platform and services.

70% More cost effective
than leading cloud
providers

40% Faster than native
OpenStack in LiINPACK
Performance tests

6x Performance boost
thanks to high
performance vNICs

FLEXIBLE INFRASTRUCTURE

Deploy what you need when you need it, mix & match different technologies and applications

SCALE-OUT STORAGE

Featuring a parallel file system for high performance IO.

UP TO 200 VMS PER APPLIANCE

Run up to 200VMs in a single Hyper-Converged 2U appliance

VSCALER

A Private Cloud Appliance
for any workload



AI & Deep
Learning



Scientific
Research



Broadcast
& Media



Automotive

USE-CASES

vScaler finely tuned cluster management software enables you to deploy clusters or clouds of any scale either on premise, in the cloud or across both in a hybrid model. Build and configure:

- HPC Clusters.
- Hadoop Clusters (Hortonworks or Cloudera)
- Clouds (Openstack)
- Parallel file systems (Lustre or BeeGFS)

CASE STUDY

CHALLENGE

NxAARK's primary objective was to find a cost-effective cloud model that would enable it to quickly offer Hosted Enterprise Cloud to its customers at a competitive price point.

SOLUTION

The vScaler HCI modular solution delivered a fully integrated, multi-location platform, hosted on-premise. The flexibility & fast learning curve aligned with NxAARK's build as-you-grow approach & enabled them to 'Go-Live' within weeks.

ADVANTAGE

By deploying an out of the box, repeatable solution, NxAARK saved on the cost of setting up a dedicated team to build a cloud platform from scratch as well as time & effort to re-invent technology upgrades & enhancements.

HYPERSCALE STORAGE PROCESSING

GPU ALTERNATIVE TO RAID



The Boston Igloo Nebari, developed in partnership with NYRIAD features NSULATE, a Linux block device that functions as a software-defined alternative to RAID for configuring fast, reliable, larger scale storage solutions.

This solution utilises modern GPUs to perform storage-processing operations and can be configured as an alternative to RAID in the same environments that RAID solutions are typically used, in addition to enabling many new processing solutions.

NYRIAD



KEY FEATURES

NSULATE

- Enables real-time hyperscale erasure coding up to 255 parity
- Cryptographic checksums and real-time corruption recovery
- Create highly parallel arrays with hundreds of devices
- High performance, even with massive degradation
- Compatible with all Linux filesystems and applications

STORAGE DENSITY

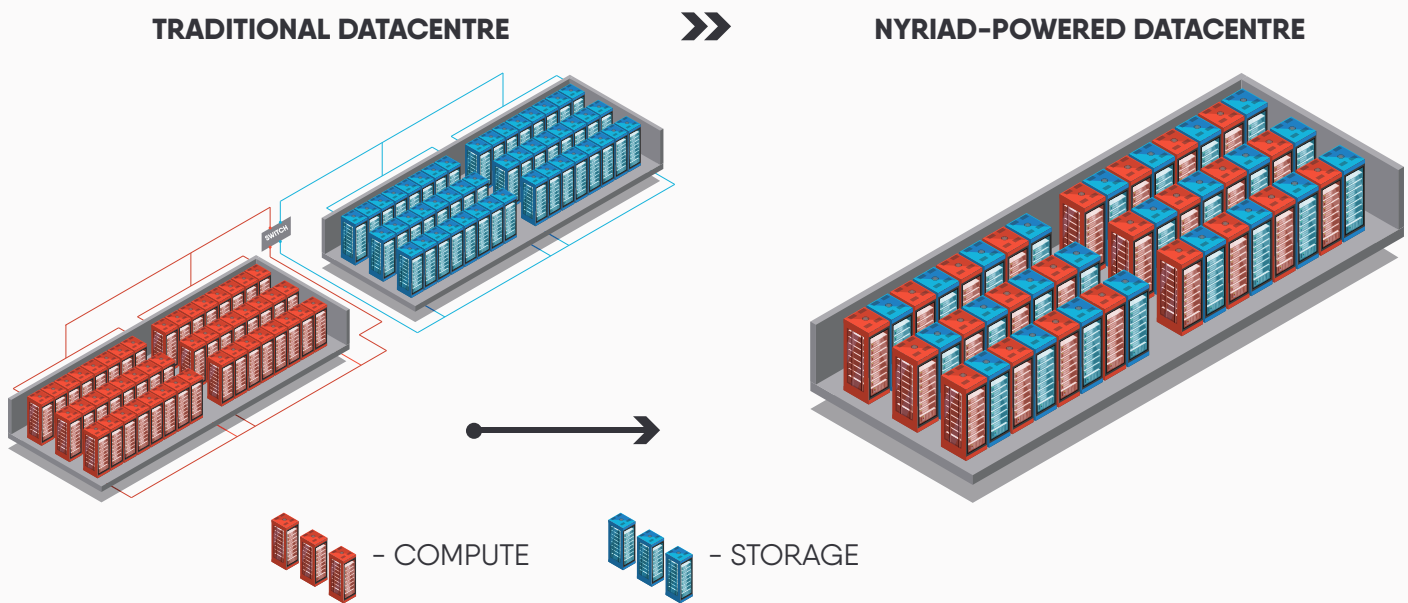


PERFORMANCE



ENERGY USAGE





USE-CASES



HPC



CLOUD



BIG DATA

WHY WOULD YOU USE A VIDEO CARD TO RUN A FILE SYSTEM?

CHALLENGE

The rapid advances in GPU computing performance has presented new challenges for storage technology. It is increasingly difficult to move data to and from storage fast enough to keep up with the processing speed of GPUs.

SOLUTION

NSULATE had to meet several requirements to be a practical solution for hyperscale storage. Including being highly resilient to failure, extremely high in performance and scalability and provably able to detect and recover from nearly all data corruption.

ADVANTAGE

NSULATE excels at real-time, high parity erasure coding, adding these capabilities automatically to file systems that do not have these features.

SCALE-OUT

PARALLEL STORAGE SOLUTION



BeeGFS® is a pure software solution for scale-out parallel network-accessible storage, developed with a strong focus on performance and designed for very easy installation and management.

The flexibility, robustness, and outstanding performance of BeeGFS allows customers to increase productivity by delivering results faster and by enabling new data analysis methods that were not possible without the advantages of BeeGFS.



KEY FEATURES

- Highly scalable parallel storage
- Uses a native Linux kernel module, so no kernel patching required
- Storage Pools enable transparent storage tiering
- Optional graphical interface for administration and monitoring systems
- Built in support for data replication through Buddy Mirroring to maximise data availability.


PROVEN BY CUSTOMERS WORLD-WIDE:

BeeGFS® is used all around the globe to provide extremely fast access to storage systems of all kinds and sizes, from small scale up to enterprise-class systems with thousands of hosts and including some of the fastest supercomputers in the world.

BENCHMARKS


METADATA OPERATIONS

BeeGFS® was designed for extreme scalability. In a testbed with 20 servers and up to 640 client processes (32x the number of metadata servers), BeeGFS® delivers a sustained file creation rate of more than 500,000 creates per second, making it possible to create one billion files in as little time as about 30 minutes.



THROUGHPUT SCALABILITY

In the same testbed system with 20 servers - each equipped with a single node local performance of 1332 MB/s (write) and 1317 MB/s (read) - and 160 client processes, BeeGFS® a sustained throughput of 25 GB/s - which is 94.7 percent of the maximum theoretical local write and 94.1 percent of the maximum theoretical local read throughput.



USE-CASES

Create Parallel File Systems on demand and on the fly: BeeGFS® on-demand allows the creation of a BeeGFS® on a set of nodes with one single command line. Possible use cases for the tool are manyfold, a few include setting up a dedicated parallel file system for a cluster job across the compute nodes used for the job, cloud computing or for fast and easy temporary setups for testing purposes.

CASE STUDY

CHALLENGE	SOLUTION	ADVANTAGE
BeeGFS® needed to develop and deliver a high performance system with a very high fault tolerance.	The easy-to-install, parallel filesystem was perfect for the I/O intensive workloads required. As a result, the customer was able to transparently spread user data across multiple servers.	The customer benefited from the holistic approach of this fully customised turnkey solution.

SCALE-OUT

OBJECT STORAGE SOLUTION



90% of the data in the world has been generated in the last 2 years. Enterprises faced with designing storage infrastructure to handle this explosive growth can no longer rely on the scalability and accessibility of legacy file system technology. They need universal access, space efficiency, and highly available, robust architectures with the freedom

to upgrade hardware as new technology becomes available. Caringo Swarm™ provides massively scalable, self-managing storage that unifies data silos and simultaneously handles mixed-use cases with a single deployment so you can extract the value of your data.

caringo



KEY FEATURES

- Built-in features to protect from accidental deletion
- 100% availability with no downtime upgrades
- Guaranteed data-integrity while continuously evolving hardware
- Choice of hardware, not locked into one vendor
- Single or multi-site deployment
- Simple, secure filetransfer with no storage silos

CARINGO SWARM BY THE NUMBERS

95%

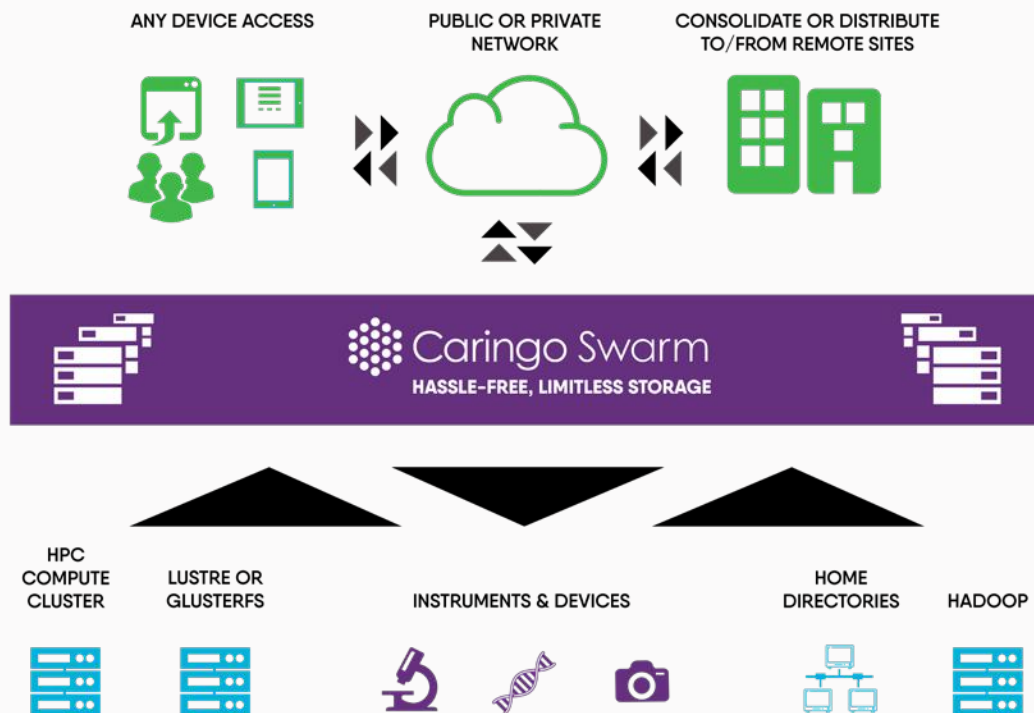
USABLE DISK
CAPACITY

70%

REDUCTION IN
POWER USAGE

75%

REDUCTION
IN TCO



USE-CASES



LOWER STORAGE TCO

- Reduces CAPEX & OPEX
- Reduces strain on primary storage



ELIMINATES RISK

- 100% business continuity
- Enables regulatory compliance



IMPROVES PRODUCTIVITY

- Manages & organises data at scale
- Unifies data silos

CASE STUDY

CHALLENGE

A university needed an object storage platform; open source solutions were discarded because of their potential complexity in system management & tuning, while traditional storage systems were not cost competitive & lacked the flexibility needed.

SOLUTION

Today, with the right technology, object storage can be adopted by organisations of all sizes. Caringo does exactly that, it has the right solution for allowing IT organisations to start with a small initial investment & grow as data & applications require.

ADVANTAGE

Chosen for the ease of use and the affordable price point in both hardware, storage and networking as well as the software; alongside the expertise that came as part of the project.

SCALE-OUT

FILE AND OBJECT STORAGE SOLUTION



Quobyte facilitates the entire HPC workflow no more silos, no more tedious and complex capacity planning, and better economics thanks to operational efficiency at scale.

Quobyte is a software defined storage solution with the capability to create multiple volumes,

providing both file and object storage. This makes Quobyte a highly versatile file system that can be utilised in several diverse scenarios, ranging from a standard Linux NFS server to S3 storage for cloud based solutions.



KEY FEATURES

- Facilitates the entire HPC workflow – no more silos, no more tedious and complex capacity planning, and better economics thanks to operational efficiency at scale.
- High-performance – the distributed parallel file system delivers all the power you need
- Massively scalable – start with a few drives and scale linearly to hundreds of PBs and beyond
- All interfaces – access data through native clients for Linux, Windows, and macOS or use S3, NFS, SMB, and Hadoop

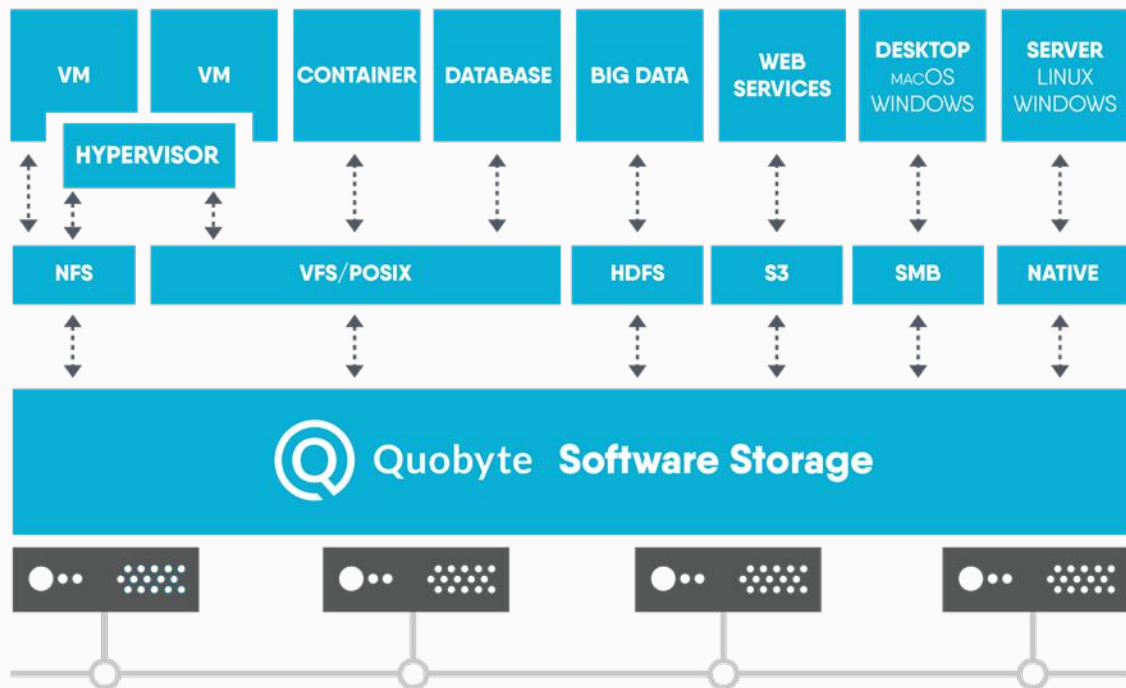
QUOBYTE

File, Block & Object Workloads within a Single System

Distributed File System

Linear Scalability & Performance

Data Safety, Redundancy Availability, & Integrity in software



USE-CASES



LIFE SCIENCES



MEDIA & ENTERTAINMENT



FINANCIAL SERVICES



CLOUD



CONTAINER



HPC & BIG DATA

CASE STUDY

CHALLENGE

A Datacentre File System was needed to manage the JASMIN Phase 4 super-data-cluster with an initial managed capacity of 42PB (petabytes). An average of 1–3PB of data is processed every day, which is expected to expand to 300PB by 2022.

SOLUTION

Quobyte's Datacentre File System gives the JASMIN data facility the ability to unify their file, block, and object storage datasets in a centralised environment consisting of 11,500 cores on 600 nodes.

ADVANTAGE

In addition to having S3 connectivity, the Quobyte system also afforded the JASMIN administrators the ability to scale their storage capacity (currently at 42PB) & performance linearly while providing ease of management.

SCALE-OUT

NVME/BURST BUFFER SOLUTION



Solving one of the biggest challenges in computing today, Excelero enables low latency & high-performance access to NVMe volumes, delivered to remote clients with minimum processor overhead. This solution is available in disaggregated or converged configurations, and employs Mellanox technologies to keep processor overhead

and latency low, whilst bandwidth is plentiful, making it a perfect fit for a myriad of different use cases. NVMesh is a Software-Defined Block Storage solution that features Elastic NVMe, a distributed block layer that allows unmodified applications to utilise pooled NVMe storage devices across a network at local speeds and latencies.

Excelero



KEY FEATURES

NVMESH

- 100% Server SAN
- Elastic NVME pools storage across a network at local speeds and latencies
- 0% CPU enables 100% converged infrastructure
- Virtual array
- Scalable
- Flexible
- Efficient
- Easy to use, manage & monitor
- Optimised for performance

SCALE & PERFORMANCE


Local performance accross the network. Predictable application performance. Smart insights in utilisation.

EFFICIENCY

Maximise the utilisation of your flash media. Reduce your capacity overhead. Easily manage & monitor.

FLEXIBILITY

Utilise any hardware. Use existing network infrastructure. Choose from multiple redundancy options.



USE-CASES



**HIGH-FRAMERATE
STORAGE FOR
ANY-K VIDEO POST
PRODUCTION**



**LOCAL BURST
BUFFER**



DATABASES



CONTAINERS

CASE STUDY

CHALLENGE

InstaDeep™ decided to build an “AI as a Service” offering. Their key requirements for the data centre infrastructure were that it needed to be scalable, flexible and highly efficient to deliver high ROI.

SOLUTION

InstaDeep chose Excelero's NVMesh® on Boston Flash-IO Talyn storage to provide GPUs with access to a scalable pool of high-performance NVMe to ensure full utilisation of the GPU processing power.

ADVANTAGE

Excelero's NVMesh eliminates any compromise between performance & practicality, & allows GPU optimised servers to access scalable, high performance NVMe flash storage pools as if they were local flash, resulting in higher ROI & faster time to results.

CLUSTER MANAGEMENT

BRIGHT COMPUTING

Bright Computing automates the process of building and managing Linux clusters in your datacentre, the cloud, and at the edge. With Bright, you can:

- Provision 3 to 30,000+ nodes from bare metal in minutes
- Repurpose servers to accommodate fluctuating workloads on the fly
- Diagnose and resolve problems and performance issues quickly
- Deploy software updates across the entire cluster effortlessly
- Extend your on-premises environment to the cloud dynamically



KEY FEATURES

- Complete, Deploy, provision, monitor, manage, and scale HPC, Big Data, and OpenStack
- Elastic - Allocate compute resources dynamically, including into the cloud
- Intuitive - Easy to install, learn, and use
- Powerful - Manage complexity and ensure scalability of clusters and clouds
- Productive - Maximise throughput, minimise effort, optimise resources
- Proven - Join hundreds of customers who have already standardised on Bright

VERTICAL INDUSTRIES


Manufacturing
and Energy

Life Sciences

Financial
Services

Government

Academic and
Research



USE-CASES



HPC CLUSTERS



EDGE COMPUTING



CLUSTERS ON DEMAND



DEEP LEARNING & AI



MANAGE CONTAINERS



CLOUD BURSTING

CASE STUDY

CHALLENGE

MSU is recognised for its research prominence, and expenditures typically exceed \$100m annually. The IT Centre Research Cyberinfrastructure group was charged with building an affordable, scalable & easy to administer shared computing resource.

SOLUTION

The Boston-built cluster sits on-premise, using virtualisation and cloud computing to provide a flexible & affordable solution, using Bright Cluster Manager to reconfigure nodes quickly and reliably.

ADVANTAGE

The fully configured cluster was first soak-tested using Bright Computing at Boston Labs, this greatly reduced the installation time on-site and meant the system was ready to use almost from day 1.

COOLING

BOSTON IMMERSED COMPUTING

Boston, in partnership with Asperitas, are making sustainable data centres possible anywhere they are needed, to facilitate emerging digital technologies, from high density compute to cloud. Immersed Computing® is a concept based on efficient liquid immersion technology and the basis for a unique infrastructure solution: the AIC24.

The AIC24 is the only immersion plug & play solution based on natural convection. Immersed Computing® is enabling optimised and even customised server designs for record breaking high density platforms (CPU and GPU) and solutions for HPC without the datacentre complexity.



KEY FEATURES

SUSTAINABILITY

- Circular design
- CO2 ↓ upto 40%
- 98% energy reuse

FLEXIBILITY

- Modular system
- Plug & Play
- Anywhere

EFFICIENCY


- TCO ↓ upto 40%
- Availability ↑
- Density ↑ x5

INDUSTRY VALUE CHAIN EXAMPLE


CLOUD VALUE CHAIN	SINGLE TENANT	HOSTING (IAAS/PAAS)	PRIVATE IN COLOCATION	SAAS IN COLOCATION
ENERGY EFFICIENCY				
SOFTWARE EFFICIENCY				
IT HARDWARE				
DC OPERATIONS				
DC FACILITIES				
DC BUILD				



1. HIGH DENSITY SERVER SOLUTIONS OPTIMISED FOR IMMERSION



2. UNIQUE IMMERSED COMPUTING SOLUTION BASED ON NATURAL CONVECTION



3. SIMPLIFIED MODULAR HIGH DENSITY DATA CENTRE DESIGN

USE-CASES

The opportunities for immersion computing in the datacentre are endless; with deployments already in HPC, Cloud, Enterprise, Life Sciences and Media & Entertainment – more and more businesses are recognising the value in what immersed computing offers.

CASE STUDY

CHALLENGE

The Department for Applied Bioinformatics in conjunction with the LOEWE Centre for Translational Biodiversity Genomics (LOEWE-TBG) was looking for a solution, which was highly energy efficient, sustainable and reliable.

SOLUTION

With the latest AMD EPYC™ processors, the Asperitas Immersed Computing® solution does not require any raised floor, CRAC units or Chillers, drastically saving customer's CAPEX and OPEX.

ADVANTAGE

Cost effective solution, able to support latest technologies & extremely low maintenance cost. Ability to extend and upgrade with only little effort.

“LEADING-EDGE DISRUPTIVE TECHNOLOGY INTEGRATORS”



vSCALER

PRIVATE CLOUD



FLASH-IO TALYN

BURST BUFFER SOLUTION



DGX-2™

GPU ACCELERATED SOLUTION



SPEAK TO US ABOUT THE **NVIDIA® DEEP LEARNING INSTITUTE**

AMD
Partner Program
ELITE



intel
Technology
Provider
Platinum 2019

WHY CHOOSE US?

TAILOR-MADE SOLUTIONS

Boston has the knowledge and expertise to tailor your ideal solution.

BOSTON HPC LABS

Remotely test and benchmark your technologies.

LEADING-EDGE TECHNOLOGY

Boston's R&D Labs facility offer the latest technology first.



STORAGE



WORKSTATIONS



SERVERS



NETWORKING



SOLUTIONS



CLOUD SERVICES

NETWORKING



Thanks to our close partnership with Mellanox, a leading end-to-end network solution manufacturer, we have a vast array of solutions to suit your HPC requirements.

NETWORKING



SOFTWARE DEFINED NETWORKING (SDN) AND CUMULUS® LINUX®

SDN changes the fundamental way that networks work, the core concept of SDN is decoupling the two roles that networks must perform; the control plane and the data plane. The addition of Cumulus Linux means customers have the option to choose the leading Linux NOS on the market.



INFINIBAND SWITCHES AND ADAPTERS

Mellanox combine industry standard InfiniBand technology, supporting up to 200Gb/s in their adapter and switch products, with integrated InfiniBand Router and InfiniBand to Ethernet gateways, Mellanox switches provide scalable fabric for powering the world's largest and fastest high-performance computing systems and next generation datacentres.



ETHERNET SWITCHES AND ADAPTERS

Mellanox Ethernet products enable users to benefit from far more scalable, lower latency, and virtualised fabric with lower overall fabric costs and power consumption, greater efficiencies, and more simplified management than traditional Ethernet fabrics.



CABLING

The Mellanox LinkX product family of cables and transceivers provides the industry's most complete line of 10, 25, 40, 50, 100 and 200Gb/s interconnect products. They are often used to link top-of-rack switches downwards to servers, storage & appliances and upwards in switch-to-switch applications. Products are available in both Ethernet and InfiniBand protocols and SFP & QSFP form factors.



NETWORKING SOLUTION

MELLANOX QUANTUM™ HDR SWITCH QM8700



Built with Mellanox's Quantum™ InfiniBand switch device, the QM8700 provides up to forty 200Gb/s ports, with full bi-directional bandwidth per port. The QM8700 is the world's smartest network switch, designed to enable in-network computing through the Co-Design SHARP (Scalable Hierarchical Aggregation and Reduction Protocol) technology. The Quantum™ switch improves the performance of selected collective operations by processing the data as it traverses the network, eliminating the need to send data multiple times between end-points.



KEY FEATURES:

PERFORMANCE

- 40 X HDR 200Gb/s ports in a 1U switch
- 80 X HDR100 100Gb/s ports (using splitter cables)
- 16Tb/s aggregate switch throughput
- Sub-90ns switch latency

ADVANCED DESIGN

- Adaptive Routing
- Congestion Control
- Collective offloads (SHARP)
- VL mapping (VL2VL)

OPTIMISED DESIGN

- 1+1 Redundant & hot-swappable power
- N+1 Redundant & hot-swappable fans
- 80 Plus Gold and Energy Star certified power supplies
- x86 ComEx Broadwell CPU

INTERESTED?

The QM8700 together with Mellanox ConnectX®-6 adapter card support HDR100. By utilising two pairs of two lanes per port, the QM8700 can support up to 80 ports of 100G to create the densest TOR switch available in the market.

NETWORKING SOLUTION

MELLANOX SPECTRUM SWITCH SN2010



The SN2010 switch is the ideal top of rack (ToR) solution for hyper-converged and storage deployments. Packed with 18 ports of 10/25GbE and 4 splittable ports of 40/100GbE, the SN2010 can deliver up to 1.7Tb/s aggregate throughput.

With its optimisation for RoCE, full buffer utilisation, and zero packet loss combined into a small form factor, the SN2010 is driving forward the world's most innovative datacentre infrastructures.



KEY FEATURES:

THROUGHPUT

1.7Tb/s
2.52B packets-per-second

LOWEST POWER

57W (ATIS)
power consumption

HALF-WIDTH FORM FACTOR

LOWEST LATENCY

300nsec for 100GbE port-to-port
Consistently low latency regardless of packet size, or L2 vs L3 forwarding
Side by side configuration
Small port count

INTERESTED?

Boston are official distributors of Mellanox and are fantastically placed to advise you on the best solution for your requirements.

Our team are ready to answer any questions you may have about networking.

HPC NODES



COMPUTE



6029TP-HTR (PER NODE)

Form Factor: 2U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
Memory: 16x DDR-4 Registered ECC 2666MHz
Expansion: 2x PCIe 3.0 x16 LP Slots
Disk: 3x 3.5" Hot swap SATA3 bays



2029TP-HTR (PER NODE)

Form Factor: 2U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
Memory: 16x DDR-4 Registered ECC 2666MHz
Expansion: 2x PCIe 3.0 x16 LP Slots
Disk: 6x 2.5" Hot swap SATA3 bays



1029TP-DTR (PER NODE)

Form Factor: 1U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
Memory: 16x DDR-4 Registered ECC 2666MHz
Expansion: 2x PCIe 3.0 x16 LP Slots
Disk: 4x 2.5" Hot swap SATA3 bays

GPU



ANNA

Form Factor: 1U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
GPU: 4x NVIDIA® Tesla® V100 SXM2
Memory: 12x DDR-4 Registered ECC 2666MHz
Disk: 4x 2.5" SATA/SAS bays



ANNA XL

Form Factor: 4U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
GPU: 8x NVIDIA® Tesla® V100 PCIe
Memory: 24x DDR-4 Registered ECC 2666MHz
Disk: 24x 3.5" Hot swap SATA/SAS bays



ANNA XL X16

Form Factor: 10U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
GPU: 16x NVIDIA® Tesla® V100 SXM2
Memory: 24x DDR-4 Registered ECC 2933MHz
Disk: 6x 2.5" SATA drives, 16x U.2 NVMe SSD
 2x M.2 NVMe drives

STORAGE



1029P-NR32R

Form Factor: 1U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
Memory: 24x DDR-4 Registered ECC 2666MHz
Network: 2x 10GBase-T LAN ports via Intel® X540
Disk: 32x Hot swap NVMe Ruler SSDs



1029UZ-TN20R25M

Form Factor: 1U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
Memory: 24x DDR-4 Registered ECC 2666MHz
Network: 2x 25GbE SFP28 ports
Disk: 20x 2.5" Hot swap drive bays



6029P-E1CR12H

Form Factor: 2U Rackmount
CPU: Dual Intel® Xeon® Scalable Family
Memory: 16x DDR-4 Registered ECC 2666MHz
Network: 2x 10GBase-T LAN ports with Intel X722 + PHY Intel X557
Disk: 12x 3.5" Hot swap SATA/SAS bays

HPC NODES



AMD EPYC™



1123US-TN10RT

Form Factor:	1U Rackmount
CPU:	Dual AMD EPYC™ 7000-series
Memory:	32x DDR-4 Registered ECC 2666MHz
Network:	2x 10GBase-T LAN ports via Intel® X540
Disk:	10x 2.5" Hot swap U.2 NVMe bays



2123US-TN24R25M

Form Factor:	2U Rackmount
CPU:	Dual AMD EPYC™ 7000-series
Memory:	32x DDR-4 Registered ECC 2666MHz
Network:	2x 25Gb SFP28 LAN Ports
Disk:	24x 2.5" Hot swap U.2 NVMe drive bays



2123BT-HNR (PER NODE)

Form Factor:	2U Rackmount
CPU:	Dual AMD EPYC™ 7000-series
Memory:	16x DDR-4 Registered ECC 2666MHz
Expansion:	2x PCIe 3.0 x16 LP Slots
Disk:	6x 2.5" Hot swap NVMe bays

GPU



DGX STATION™

Form Factor:	Tower
CPU:	Dual Intel® Xeon® E5-2698 v4
GPU:	4x NVIDIA® Tesla® V100 SXM2
Memory:	256GB DDR-4 Registered ECC
Disk:	3x 1.92TB SSD RAID 0



DGX-1™

Form Factor:	3U Rackmount
CPU:	Dual Intel® Xeon® E5-2698 v4
GPU:	8x NVIDIA® Tesla® V100 SXM2
Memory:	512GB DDR-4 Registered ECC
Disk:	4x 1.92TB SSD RAID 0



DGX-2™

Form Factor:	10U Rackmount
CPU:	Dual Intel® Xeon® Platinum 8168
GPU:	16x NVIDIA® Tesla® V100 PCIe
Memory:	1.5TB DDR-4 Registered ECC
Disk:	8x 3.84TB NVMe SSD

STORAGE



6049P-E1CR36H

Form Factor:	4U Rackmount
CPU:	Dual Intel® Xeon® Scalable Family
Memory:	16x DDR-4 Registered ECC 2666MHz
Network:	2x 10GBase-T LAN ports with Intel X722 + PHY Intel X557
Disk:	36x 3.5" Hot swap SATA/SAS bays



6049P-E1CR60H

Form Factor:	4U Rackmount
CPU:	Dual Intel® Xeon® Scalable Family
Memory:	24x DDR-4 Registered ECC 2666MHz
Network:	2-ports of 25GbE SFP2
Disk:	60x 3.5" Hot swap SATA/SAS bays



6048R-E1CR90L

Form Factor:	4U Rackmount
CPU:	Dual Intel® Xeon® E5-2600
Memory:	8x DDR-4 Registered ECC 2400MHz
Network:	Intel® i350 Dual Port Gigabit Ethernet
Disk:	90x 3.5" Hot swap SATA/SAS

TECH TALK

KNOW YOUR LIMITS | AERODYNAMICS IN MOTORSPORTS

INTRODUCTION

On January 1st 2018 the FIA introduced the latest evolution of the Aerodynamic Testing Restrictions for Formula 1 and with it came the biggest change in the CFD restrictions since they were first introduced back in 2009. In this article we go behind the scenes with the HPC specialists Boston Ltd to look at the impact of these changes and how the F1 teams have benchmarked new solutions and upgraded their CFD supercomputers.

BACKGROUND – THE INTRODUCTION OF F1 AERO RESTRICTIONS

To put these latest changes into context we need to understand the history of the restrictions both for CFD and the wind tunnel. In 2008 aerodynamic testing was at its peak. BMW Sauber, Honda, Williams and Toyota had all invested huge sums of money in new state of the art full size wind tunnels, each costing tens of millions of pounds. All the top teams were operating in two wind tunnels simultaneously, and for Toyota they were utilising not one but two full size wind tunnels 24/7. Of course the vast majority of this wind tunnel testing was being undertaken with scale models, and over the years the scale of these models had increased from 40% to 50% and then up to 60%. Operating in two wind tunnels 24/7 would allow these teams to undertake around 500 wind tunnel simulations per week, with each simulation incorporating approximately 20 different car attitudes. Full size wind tunnel testing was also commonplace with teams either using their own facility or using a customer facility such as Windshear in the USA. In 2008, teams were already using CFD routinely as part of the aerodynamic development process, but

as the software and correlation improved and the cost of the hardware reduced so the teams were starting to use it more and integrate it further into the design cycle. BMW Sauber were leading the way in CFD hardware at that time with the Albert 3 supercomputer and over 4000 Intel cores, but other leading teams were not far behind. It was clear that something had to be done to curb the growth of aerodynamic testing in F1 and the costs associated with it. The first step came into force on January 1st 2009 as part of the FOTA Resource Restriction Agreement (RRA). This controlled the aerodynamic resources the teams could deploy via restrictions on the wind tunnel “wind on time” (WON) and the CFD compute capacity measured in TFLOPS. Wind On Time was simply a measure of the amount of time the fan was turned on in the wind tunnel and the wind speed in the test section was above 15m/s. For CFD, the capacity was measured in TFLOPS - effectively the number of floating point operations completed within the designated 8 week accounting period for the regulations, as defined by the following equation:

$$\text{TotFLOPs} = \left(\frac{\text{MFPPC} \times \text{CCF} \times \text{NCU} \times \text{NSS}}{604,800 \times 8 \times 1,000} \right)$$

Where:

TotFLOPs = total number of TeraFLOPs used per second

MFPPC = Peak double precision floating point operations per cycle per core of the processing unit

CCF = Peak processing unit clock frequency in GigaHertz

NCU = Number of processing unit cores used for the run

NSS = Number of solver wall clock seconds elapsed during the run

TECH TALK

2018 CFD REGULATION CHANGES

By 2012 AMD had been persuaded by many teams to produce a limited run of Fangio chips allowing more of the grid to upgrade their supercomputers to this specification with most of the remaining teams running an Intel® Ivybridge system. With the FiA not willing to extend the flops/cycle exemption to more modern Intel® chips, such as the V3 Haswell CPUs which were rated at 16 dp flops/cycle, and AMD not producing any more Fangio chips, the teams were now locked into these older systems purely by virtue of the regulations. Newer chips were simply not viable because of their high flops/cycle rating. These older systems were coming to the end of their life and were no longer supported by Intel® or AMD. Clearly the FiA had to do something, and the target was to introduce a new regulation which aligned the F1 aero departments more closely with the wider CFD industry and allowed the teams to upgrade to more modern, supported technology. This resulted in the 2018 CFD restrictions and a move from TFLOPS to Mega Allocation Unit Hour.



BENCHMARKING 2018 REGULATIONS WITH BOSTON LTD

Boston Ltd have been specialising in High Performance Computing in a wide range of sectors for over 25 years and in 2017 they formed a new partnership with Tim Milne of UniFi Engineering Services Ltd (UniFi) and Dr. Lee Axon of Computational Engineering Ltd (CE).

Tim and Lee have extensive F1 experience, most recently at Manor F1 where they were Head of Aerodynamics and Head of CFD Correlation respectively. This group combined Boston's extensive HPC technical knowledge with UniFi's and CE's F1 aerodynamics and CFD experience to provide the F1 teams with a comprehensive benchmarking of the new AMD EPYC™ and Intel® Skylake Platforms. They were able to use all the main F1 CFD codes with models aligned to F1 methodologies and HPC hardware setups to extract the maximum possible performance from the new regulatory environment.

The project began in August 2017 when Boston Ltd were one of the first companies worldwide to have invested in their own 8 node dual socket AMD EPYC™ system based on the EPYC7601 32 core chips and a similar 8 node system based on the Intel® Skylake 8176 Platinum 28 core chip. Following the benchmarking of the older Ivybridge system, a number of options within the AMD EPYC™ range as well as the Skylake 8176 chip were evaluated as single node tests to gain an initial assessment of the various different chips available in each family, as well as some insight into the time/iteration performance benefits of different options such as the turbo mode.

Once testing migrated onto the full, multi-node systems the full optimisation process could begin. This involves running the same model over a wide range of different setups, including options for memory bandwidth per core used and process bindings. Finally, Boston were able to use their close links with AMD to further optimise for the requirements of the F1 environment increasing the memory bandwidth whilst retaining a relatively low base clock speed.

TECH TALK



CONCLUSIONS – HAS THE REGULATION CHANGE ACHIEVED THE DESIRED RESULT?

The stated targets of the FiA for this change in regulations was to enable the F1 teams to upgrade from their Fangio & Ivybridge systems to the latest technology available, but without a clear performance pressure to do so, and with the aim of better aligning the F1 industry with the wider CFD industry. It is against these targets that the effectiveness of the regulatory change should be judged.



Firstly it is clear that all the F1 teams have upgraded to a new system, with most teams having done so ahead of the regulatory change date of 1st January 2018. The first aim has been achieved – the Fangio and Ivybridge systems that the teams were operating have now been obsoleted. However the benchmarking work completed by Boston clearly demonstrates the huge performance advantage available by purchasing a new multi-million pound system. This was not the aim of the new regulations, in addition Boston suggests that the teams are now able to complete between 1000 and 1500 CFD simulations per week based on a typical CFD model of around 200m cells. The teams may elect to “trade” some of this capacity for larger models (some teams are known to run CFD models of approaching 1bn cells) or better quality models (transient simulations rather than steady state).

Typically each generation of compute chip that is released by AMD/Intel® delivers 10 – 20% improvement in efficiency. Back in 2009 this would give the teams an extra 10 – 20 CFD runs per week, and therefore would not easily justify the large cost in replacing their CFD clusters. In 2018, with the massive increase in capacity, the same 10 – 20% improvement available from each evolution of compute chip technology is 100 – 200 runs – that is the same as the total capacity of the systems in 2009.

Is this a bad thing? Arguably not. HPC systems now are much cheaper now than they were in 2009. The FiA focus remains on reducing wind tunnel reliance and delivering greater CFD capacity in exchange, and the current regulations deliver that.

However does it help to “level the playing field” between the high budget teams and the low budget teams? Does it help to encourage new teams into the sport? And does it make the working practices within the F1 aero departments more aligned to the wider CFD industry? With AMD releasing their second generation of EPYC™ chip in 2019, the reaction of the teams will be interesting. Will they all upgrade immediately? Or will the well funded teams take the opportunity to get a performance advantage from the new technology that the smaller teams cannot afford? Time will tell...

By Gemma Hatton

as featured in Racecar Engineering Magazine

Boston, UniFi and CE continue to develop their partnership with a focus on the F1, motorsport and automotive industries across all CFD codes and working practices. For further information go to www.boston.co.uk, www.unifimotorsport.com or www.computationalengineering.co.uk

TECH TALK

INSTADEEP™ POWERS AI AS A SERVICE WITH SHARED NVME

NVMESH™ FEEDS UNLIMITED STREAMS OF DATA TO GPU-BASED SYSTEMS WITH LOCAL PERFORMANCE


InstaDeep™ Ltd. is a global AI innovator, headquartered in London with offices in Paris, Tunis, Nairobi and Lagos. The company delivers AI-powered decision-making systems for customers across a wide array of industries, including Logistics, Manufacturing, Energy and Mobility. With expertise in both machine intelligence research and concrete business deployments, InstaDeep™ provides competitive advantages to customers in an AI-first world: the company's AI solutions allow enterprises to unlock data insights, realise value, increasing efficiency and speed across organisations.

UNCAGING GPUS AND NVME PERFORMANCE FOR AI AND MACHINE LEARNING

AI and ML use has exploded over the past few years as four key technology evolutions have made it far easier to capture, store and process data into insights that can help enterprises outsmart the competition:

- 1) New sensor technologies have proliferated that capture images, temperature, heartrate, and more – adding even more data volumes.
- 2) Big Data analytics and Data Lakes for storing these massive volumes of data arose, so that teams could analyse and mine that data to turn it into valuable business or research insight.
- 3) The rise of powerful GPU technologies that lower the cost of massive compute on those data have made parallel processing faster and much more powerful.

4) Next-gen storage options such as NVMe flash media have swept the storage industry and are well-suited to these new computational engines, although they harken back in time to the days when direct attached storage (DAS) models were new. DAS is fast, but often underutilised and hence costly.



INSTADEEP™ AI AS A SERVICE

When InstaDeep™ decided to build an AI as a Service offering, it had a few key requirements in mind. The data centre infrastructure needed to scale modularly, as the company began this cutting-edge service with a few key customers but planned to quickly expand to its entire global customer base. The infrastructure had to be flexible to meet performance requirements for a wide range of workloads. Today the infrastructure is used by multiple scientists who run workloads for many different customers. Finally, for the service to be attractive for customers, and a sound business move for InstaDeep, the infrastructure needed to be highly efficient – the GPUs especially would need to deliver the highest return on investment (ROI).

InstaDeep chose the Boston Flash-IO Talyn storage with Excelero's NVMesh® to provide GPUs with access to a scalable pool of high-performance NVMe to ensure full utilisation of

TECH TALK

the GPU processing power. The Talyn system included a 2U Boston Flash-IO Talyn server with Micron NVMe flash and Excelero NVMesh software that provides access to up to 100TB external high-performance storage. Leveraging the Mellanox 100GB Infiniband network cards in the DGX, the GPUs use the NVMe storage with local performance. The ability to choose any file system to run on NVMesh was an immense benefit. Early tests quickly showed that external NVMe storage with Excelero gives equal or better performance than local cache in the DGX.



FEEDING THE GPU BEAST

The biggest advantage of modern GPU computing is also creating its biggest challenge: GPUs have an amazing appetite for data. Current GPUs can process up to 16GB of data per second. The latest NVIDIA® DGX-2™ system has as many as 16 GPUs, but by far not enough local storage. The DGX-1 has a theoretical limit of 7.8GB/s bandwidth, but with only 4 SATA SSDs it is limited to about 2.2GB/s. Theoretically it can process 2 million random IOPs but local storage only provides 400K IOPs. The latest NVIDIA® DGX-2™ has 30TB (8 x 3.84TB) local NVMe but is not optimised to use it efficiently. Other brand GPU servers typically feature few PCIe lanes for local flash (NVMe or other), meaning even the lowest latency option for these servers is a

severe bottleneck or is simply too little capacity for the GPUs. Starving the GPUs with slow storage or wasting time copying data wastes expensive GPU resources and affects ROI.

OPTIMAL APPROACH FOR SPECIFIC AI, ML USES

NVMe flash offers great benefits for specific AI use cases like training a machine learning model, and checkpoints. Machine learning involves two phases - training a model based on what is learned from the dataset and running the model. Training of a model is the most resource hungry stage. Hardware used for this phase, incorporating high-end GPUs or specialised system-on-chips (SoCs), is expensive to buy and operate so it should be always busy for best ROI.


SUMMARY

The capability of GPUs and the rise of affordable compute power challenges IT teams to think at data centre scale – leveraging the ability to apply AI, ML and deep learning techniques to large data pools, while making sure the entire system is scale-out, highly performant, and efficient. The only storage that is fast enough to keep up with GPUs is local NVMe flash, since GPUs, networking and NVMe are all competing for valuable PCIe connectivity, one of them must compromise and settle for less. The Boston Flash-IO Talyn with Excelero NVMesh eliminates any compromise between performance and practicality, and allows GPU optimised servers to access scalable, high performance NVMe flash storage pools as if they were local flash. This technique ensures efficient use of both the GPUs themselves and the associated NVMe flash. The result is higher ROI, easier workflow management and faster time to results.

TECH TALK

WHY WON'T YOU CHANGE? – THE BATTLE OF THE DATA CENTRE

The implementation and operation of all technology is the responsibility of the Head of IT Infrastructure. They are the decision maker and oversee many things such as data centre, network and server services, user support/help desk, workstation management, storage software and much more. The most common strategies the Head of IT Infrastructure would have in place in order to help achieve their objectives can include; hiring appropriate IT staff as well as training them to enhance their skill set, allocating set budgets for new technology, and positioning IT infrastructure as a strategic asset to the business, and understanding how to make it a profit centre. Enabling collaboration across different departments objectives can also be a good strategy.



There are also, of course, both internal and external challenges that could hinder the Head of IT Infrastructure from achieving their goals such as:

INTERNAL

- Turnover of key staff to other companies, resulting in them taking their knowledge elsewhere
- Lack of alignment between visionaries in the company (C-suite not aligned)
- Too much time spent on maintenance
- No internal champion of product helping to drive change

EXTERNAL

- Changes/updates or modernisation of applications from software providers
- Integration with public cloud
- Enforcement of new technologies in order to keep up with competition
- Uncertainty of how to implement new technology (Big Data, IOT, AI, Software-defined tech, NVMe)

There can also be instances where even though the Head of IT Infrastructure is not completely happy with the current status quo, they would still stick to it. These change inhibitors include no internal champions to support their vision, lack of budget, the current network is fine, resources are stretched or the slow adoption of new technologies.

Whilst the Infrastructure Manager must be able to use their knowledge to select the right technology that enables the organisation to move forward, rather than waiting on the IT department, they most likely have uncertainty of how to implement new technology, and/or lack of budget to drive change as C-suite (CEO, CTO, etc) doesn't see it being relevant.

WHAT IS LINUX?

For a modern data centre to be agile and scalable, it needs to be able to keep up with the amount of data that floods it each day, hour, minute and second. Due to this, many people turn to Linux but what exactly is Linux?

Linux is an open source operating system and can be installed on a variety of different types of hardware. The operating system effectively manages the communication between your software and your hardware so without

TECH TALK

the system, the software wouldn't function. Although Linux can be seen as similar to other operating systems, it's source code is available to the public to view, edit, and for those with the appropriate skills, to contribute to. Linux's package management system allows users to download all kinds of applications including software important for cloud infrastructure, enabling you to better customise and optimise your data centre.



DIFFERENCES BETWEEN TRADITIONAL VS. LINUX NETWORKING

DISAGGREGATION

As Linux is an open solution, it allows your network to be more customisable by choosing the hardware, ASIC and applications that you want. Having Linux as your operating system enables you to build a custom network that is suited to your unique needs.

COMMUNITY ORIENTED

As opposed to relying on a vendor every time a problem with your network occurs, Linux is monitored and maintained by members of the Linux community who are able to take matters into their own hands.

COMPLETE STANDARDISATION

To make it simple, Linux is a common language throughout the data centre meaning reduced

complexities and complete interoperability. This enables rapid availability of applications to the end user which results in management becoming easier.

Boston are pleased to offer testing with Cumulus Linux at our world class lab facilities. You will have access to the latest technologies and will have the opportunity to test them out on your own workloads. Our Boston Labs team are on hand with their expert help every step on the way. Even if you aren't able to visit our labs in person, we offer a remote dial-in service via RDP, SSH or other means which enables you to trial our solutions from wherever you are in the world!



Boston are a proud distributor of Mellanox, we are delighted to share that the entire portfolio of Spectrum based 25, 40, 50, and 100 Gb/s Ethernet switches, to bring web-scale IT efficiency to the data centre, are supported by Cumulus Linux. This partnership enables customers to automate, customise and scale using web-scale principles which can be found in the world's largest data centres.

If you would like to find out more about Open Networking or how Cumulus can benefit your data centre - get in touch with us today at sales@boston.co.uk or call us on **01727 876 100** where our team is standing by to help you!

THE FUTURE IS NOW



Asperitas AIC24 Immersed Computing®
datacenter in a box solution.
Liquid cooled high performance cloud
and compute anywhere.

sustainable plug & play high density



www.asperitas.com



BOSTON
TRAINING
ACADEMY

BOSTON
Servers | Storage | Solutions

EXPERT TRAINING FOR YOUR STAFF & CUSTOMERS IN A RANGE OF APPLICATIONS

FOR MORE INFORMATION ABOUT OUR
HANDS-ON TRAINING COURSES IN BRIGHT CLUSTER MANAGEMENT,
LINUX MANAGEMENT TOOLS, CENTOS, CARINGO, QUOBYTE AND MORE...



ALSO DELIVERING NVIDIA®
DEEP LEARNING FUNDAMENTALS



DEEP
LEARNING
INSTITUTE



Servers | Storage | Solutions



UK (Headquarters)

+44 (0) 1727 876 100

sales@boston.co.uk

www.boston.co.uk

AFRICA

+27 826 486 485

sales@boston.co.uk

www.boston.co.uk

INDIA

+91 22 5002 3262 (Mumbai office)

+91 80 4308 4000 (Bangalore office)

sales@bostonindia.in

www.bostonindia.in

GERMANY

+49 (0) 89 9090 1993

sales@boston-it.de

www.boston-it.de