

#WEAREHPC

**BOSTON**  
Servers | Storage | Solutions



**AMD**



**intel**



**IBM**

# HPC BROCHURE

**Powerful Technology**  
*Intelligently Applied™*

# CONTENTS



<b>1</b>	<b>INTRODUCTION</b>	<b>20</b>	<b>COOLING SOLUTIONS</b>
1	About Boston	20	Liquid Cooling
<b>2</b>	<b>HPC CLUSTER ARCHITECTURE</b>	21	Cooling Solutions for HPC
2	HPC Cluster Architecture & Workflow	22	ZutaCore
<b>3</b>	<b>COMPUTE SOLUTIONS</b>	24	Asperitas
3	Compute Solutions for HPC	25	Submer
4	AMD 4 <sup>th</sup> Gen EPYC™ 9004 Series Processors	<b>26</b>	<b>NETWORKING SOLUTIONS</b>
5	4 <sup>th</sup> Gen Intel® Xeon® Scalable Processors	26	Networking Solutions for HPC
<b>6</b>	<b>ACCELERATED SOLUTIONS</b>	27	NVIDIA Infiniband Quantum™-2 Platform
6	Accelerated Solutions	28	NVIDIA Ethernet Spectrum®-4 Platform
7	NVIDIA® DGX™ Family	<b>29</b>	<b>SERVICES</b>
8	AMD Instinct™ MI200 Series Platform	29	Services
9	Intel® Habana - Gaudi2	30	Boston Training Academy
10	Graphcore	32	Boston Technology Consulting
<b>11</b>	<b>STORAGE SOLUTIONS</b>	<b>33</b>	<b>BOSTON LABS</b>
11	Storage Solutions for HPC	33	Boston Labs Introduction
12	vScaler®	34	Planning, Installation, Testing & QA
14	DDN EXAScaler	35	Additional Services and Support
16	Weka IO	<b>36</b>	<b>TECH TALK</b>
18	IBM Spectrum Scale™	36	Profits through simplicity
		37	Compute Express Link™ (CXL™)

**BOSTON**  
Servers | Storage | Solutions

# INTRODUCTION

Boston Limited has been synonymous with servers, storage and HPC solutions for 30 years, establishing itself as the early adopter of a wide variety of technologies and a leading global provider of tomorrow's technology today.

Through long-standing partnerships with Supermicro and other leading innovators, it has continuously enabled innovation through collaboration, ensuring thousands of successful deployments for business transformation projects driven through digital enablement. By intelligently applying its skills with emerging technology and creating tailored customer experiences in the Boston Labs, Boston is helping every industry in its pursuit to build more cognitive systems.



**TEST DRIVE THE LATEST SOLUTIONS  
TODAY AT BOSTON LABS**





# 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**



# COMPUTE SOLUTIONS



## COMPUTE SOLUTIONS FOR HPC

Supermicro and Boston has been providing HPC solutions for the last 30 years. Much like us, Supermicro is incredibly versatile with their range of systems and their applications. Their solutions are comprised of a wide range of building blocks to meet clients' specific needs.

Reference designs already exist from Supermicro for enterprise HPC, scientific research, and AI and deep learning – all of which can be tailored to the smallest detail. If a system is certified by Supermicro, you can be sure that its components are optimised for the highest performance and the best power-efficiency.

Supermicro systems can be used in virtually any vertical. A significant portion of data centres around the world use Supermicro systems due to their reliability, density, and ease of application.

Furthermore, Supermicro also provides an extremely competitive density to price ratio. These benefits have been leveraged in CFD, FEA, weather prediction, oil and gas research, automotive, aerospace, civil engineering, hardware design, naval engineering, to name a few.

Finally, HPC requires vast amounts of compute power which in turn requires a significant amount of energy. This is so the workload is finished in a reasonable time span while minimising the power necessary to reach to goal.

Supermicro promotes greener computing by providing highly optimised and power efficient systems and configurations. This not only lowers TCO but also works toward sustainability goals.

COMPUTE SOLUTIONS

# AMD 4<sup>TH</sup> GEN EPYC™ 9004 SERIES PROCESSORS



The processor that will put the wind in your sails has been announced! First, AMD took you to Naples, then Rome, then Milan. Next stop: Genoa, the codename for AMD's 4<sup>th</sup> generation EPYC™ 9004 series CPU.

The EPYC™ family of CPUs provides great performance and security while maintaining power and cost-efficiency. AMD's 4<sup>th</sup> generation EPYC™ 9004 series CPU keeps this trend alive by doubling down on what made the previous EPYC™ CPUs so iconic.

The EPYC™ 9004 series CPU uses DDR5 memory, with 50% more channels providing 2.25 times the memory bandwidth of the previous generation.

The chip comes with twice the IO capacity and PCIe®5 making it a cornerstone of AMD's architectural leadership. All this culminates in lower TCO over three years by around 29%.

AMD



## KEY FEATURES

### EXTENDING COMPUTE LEADERSHIP


- Leadership socket and per-core performance
  - Up to 96 "Zen 4" cores in 5nm
- Leadership memory bandwidth & capacity
  - 12 channels DDR5 with up to 12TB of memory capacity



## BENEFITS


- Security features at the silicon level
- Lower TCO and energy consumption
- Excellent price-performance
- Previous AMD EPYC™ family CPUs have provided a massive upgrade for numerous businesses, and the AMD 4<sup>th</sup> generation EPYC™ 9004 series processors carries on this tradition.

AMD  
EPYC



## COMPUTE SOLUTIONS

# 4<sup>TH</sup> GEN INTEL® XEON® SCALABLE PROCESSORS



4th Gen Intel® Xeon® Scalable processors are redefining performance and features built-in accelerators to improve performance across the fastest-growing workloads in AI, analytics, networking, storage, and HPC.

Built-in accelerators can result in more efficient utilisation and power efficiency by lightning the load from CPU core resources, which also carries the benefit of working toward sustainability goals.

Other seamlessly integrated accelerators speed up data movement and compression for faster networking, boost query throughput for more responsive analytics, and offload scheduling and queue management to dynamically balance loads across multiple cores.

Zero trust security strategy is brought to life via Intel® Software Guard Extensions (Intel® SGX).

intel



## KEY FEATURES


### CPU CORE RESOURCES WITH BUILT-IN ACCELERATORS

- Intel® AMX for AI
- Intel® IAA for Data Analytics
- Intel® DSA for 5G/Networks
- Intel® DLB for Storage
- Intel® QAT for Cloud



## BENEFITS

- New standards in data center architecture
  - Multi tile SoC for scalability
  - Physically tiled, logically monolithic
  - General purpose and dedicated acceleration engines
- Designed for Cloud, Microservices & AI workloads
  - Performance core architecture
  - Workload specialized acceleration
- Pioneering advanced memory and IO transitions
  - DDR 5 & HBM
  - PCIe 5.0
  - Enhanced virtualization capabilities






# ACCELERATED SOLUTIONS




## ACCELERATED SOLUTIONS FOR HPC




### NVIDIA® DGX™ SYSTEMS

As an 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.



### INTEL® HABANA® GAUDI2®

Born for Deep Learning. Raised to a whole new level. Supermicro and Intel have partnered to bring a next-generation AI deep learning server to market.



### GRAPHCORE

Graphcore created a completely new processor, the IPU, specifically designed for AI compute. The IPU's unique architecture lets AI researchers undertake entirely new types of work to drive the next advances in machine intelligence.




### AMD

Exascale compute is rapidly becoming commonplace. Tackling workloads at this scale requires hardware to keep pace and AMD addresses this requirement with their MI series GPUs. These GPUs perform to the highest standard in HPC, AI, and machine learning.

## ACCELERATED SOLUTIONS

# NVIDIA® DGX™ H100



Cut down on training time with the newest NVIDIA® DGX™ H100. Shorter render time, faster engine compile performance, and reduced simulation times are the minimum you can expect. Harness the power of PCIe 5.0 lanes to bring out the most from your GPU and storage.

Successor to the A100 GPU, which was built on the Ampere architecture, the H100 GPU, built on the Hopper architecture, is a significant

upgrade compared to the A100. It is capable of everything that its predecessor is, and much more. With 80 Billion transistors, and 4.9 TB/s bandwidth, twenty H100 GPUs can sustain the equivalent of the entire world's internet traffic. Such bandwidth is made possible by leveraging NVLink, which is NVIDIA's very own GPU high-bandwidth interconnect.




## KEY FEATURES

### Securely accelerate workloads from enterprise to exascale

#### Up to 9X Higher AI Training on Largest Models


Mixture of Experts (395 Billion Parameters)



Projected performance subject to change. Training Mixture of Experts (MoE) Transformer Switch-XXL variant with 395B parameters on 1T token dataset | A100 cluster: HDR IB network | H100 cluster: NVLink Switch System, HDR IB


#### Up to 30X higher AI inference performance on the largest models

Megatron chatbot inference (530 billion parameters)



Projected performance subject to change. Inference on Megatron 530B parameter model chatbot for input sequence length=128, output sequence length=20 | A100 cluster: HDR IB network | H100 cluster: HDR IB network for 16 H100 configurations | 32 A100 vs 16 H100 for 1 and 1.5 sec | 16 A100 vs 8 H100 for 2 sec

#### Up to 7X higher performance for HPC applications



Projected performance subject to change. 3D FFT (4K\*3) throughput | A100 cluster: HDR IB network | H100 cluster: NVLink Switch System, HDR IB | Genome Sequencing (Smith-Waterman) | 1 A100 | 1 H100

## A TURNKEY SUPERPOD


Reach the heights of supercomputing with one simple, and all-encompassing solution: the DGX™ Superpod. Built on market-leading infrastructure, the DGX™ Superpod can consistently keep pace with other AI clusters while remaining a turnkey solution. This is achieved through leveraging highly optimised hardware to make the best use of the H100 GPUs.

## PREDECESSOR OF H100

NVIDIA® DGX™ A100 is the predecessor of the H100 GPU. It was built on the Ampere architecture. When it launched it was the largest leap in performance across all 8 generations of NVIDIA graphics hardware.

## ACCELERATED SOLUTIONS

# AMD Instinct™ MI200 Series Platform



The AMD Instinct™ MI200 series accelerators are the newest data center GPUs from AMD, designed to power discoveries in mainstream servers and supercomputers, including some of the largest exascale systems, enabling scientists to tackle our most pressing challenges from climate change to vaccine research.

The era of exascale is here. Immense computational power coupled with the fusion


of HPC and AI is enabling researchers and scientists to tackle our most pressing challenges from climate change to vaccine research.

With the AMD Instinct™ MI200 accelerators and ROCm™ 5.0 software ecosystem, innovators can tap the power of the world's most powerful HPC and AI data center GPUs to accelerate their time to science and discovery.



## KEY FEATURES

Performance	MI210	MI250
Compute Units	104 CU	208 CU
Stream Processors	6,656	13,312
Peak FP64/FP32 Vector	22.6 TFLOPS	45.3 TFLOPS
Peak FP64/FP32 Matrix	45.3 TFLOPS	90.5 TFLOPS
Peak FP16/BF16	181.0 TFLOPS	362.1 TFLOPS
Peak INT4/INT8	181.0 TOPS	362.1 TOPS



## PERFORMANCE

Based on the 2<sup>nd</sup> Gen AMD CDNA™ architecture, AMD Instinct™ MI200 accelerators deliver a quantum leap in HPC and AI performance over competitive data center GPUs today.

With an up to 4x advantage in HPC performance compared to competitive GPUs, the MI200 accelerator is the first data center GPU to deliver 383 teraflops of theoretical mixed precision FP16 performance for deep learning training, offering users a powerful platform to fuel the convergence of HPC and AI.





## ACCELERATED SOLUTIONS

# INTEL® HABANA® GAUDI2®



The original Habana® Gaudi® already boasted 40% better price performance in the AWS cloud with Amazon EC2 DL1. Habana Gaudi2, the second generation of these processors, capitalises on the success of its predecessor by focusing on high-efficiency deep learning. Habana Gaudi2 makes the already cost-efficient previous generation even more efficient by lowering training times. This level

of efficiency will increase performance in fields such as object detection in autonomous vehicles, object detection in medical imaging, and defect detection in manufacturing.



## KEY FEATURES

This level of power was attained by increasing the number of AI-customised Tensor Processor Cores from 8 to 24. The other number that increased between the two generations is the number of ports on the chip: 24 100-Gigabit RDMA over Converged Ethernet (RoCE2) ports were integrated on the chip, whereas the previous iteration had 10.

Gaudi2 enables customers to choose from a wide array of Ethernet switching and related networking equipment, enabling cost-savings – which is thanks to the widely used industry-standard Ethernet connectivity. Furthermore, the Networking Interface Controller (NIC) ports are integrated on the chip.


## SCALING SYSTEMS WITH GAUDI2

Habana has made it cost-effective and easy for customers to scale out training capacity by amplifying training bandwidth on Gaudi2 with the increased number of RoCE2 ports on chip.



## ACCELERATED SOLUTIONS

# GRAPHCORE



Graphcore, a Boston partner, and a fellow British company, specialises in semiconductors to produce accelerators for AI and machine learning. Graphcore IPU systems are already being used by Boston customers to reach new breakthroughs and take humanity to new heights. The 3rd generation BOW IPU systems by Graphcore use the revolutionary 3D Wafer-on-Wafer (WoW) processor, which, thanks to

the architecture, dramatically increases power efficiency and performance. The BOW-2000 IPU machine with its 1.4 petaFLOPS of AI performance drastically cuts down training and inferencing times. Such awesome AI power is fit into a slim 1U server blade, and with its memory communication architecture makes it easy to reach a supercomputing scale.

GRAPHCORE

## KEY FEATURES

For AI to work on its own, it must be taught and fed information, which, of course takes time. Graphcore pushed the envelope and made their BOW IPU reduce training time from days to hours by creating the world's first Wafer-on-Wafer (WoW) 3D stacked processor. It brings benefits for several verticals:

### Natural Language Processing (NLP)

- improving internet search sentiment analysis
- fraud detection
- chatbots

### Imaging


- Medical imaging
- Claims processing
- Smart cities
- Self-driving cars

### Research

- Weather forecasting
- Computational fluid dynamics
- Protein folding
- Oil / gas exploration

## GRAPHCORE BOW PODS

Graphcore's BOW Pod line harnesses the power of their IPUs to produce a flexible yet powerful AI and machine learning system. The BOW Pod16 leverages 16 BOW-2000 IPUs, providing 4 petaFLOPS of FP16 and 1 petaFLOPS of FP32 compute. Definitely boasting power to be reckoned with. Graphcore also provide this exact performance at larger scales with the BOW Pod 64, Pod 256. These set ups are also available as BOW Pod512 and Pod1024 to tackle the largest possible neural networks, AI frameworks, and models.



# STORAGE SOLUTIONS

## STORAGE SOLUTIONS FOR HPC



### HYPER-CONVERGED, VIRTUALISED & CLOUD

Boston can provision hyper-converged appliances which are built specifically for HPC workloads and are designed to consume any of the storage options below.



### VENDOR-DEFINED APPLIANCE HPC STORAGE

Appliance HPC storage provides a packaged solution which can be easily managed & tuned with the manufacturer supporting the whole stack to provide peace of mind.



### COMMERCIAL SOFTWARE DEFINED HPC STORAGE

Boston can build feature rich HPC storage on Supermicro hardware to deliver value without compromising on usability.



### LOW-COST, OPEN SOURCE, SOFTWARE-DEFINED HPC STORAGE

Boston can deliver open source software defined storage to deliver low cost per TB, ideal for enterprises with experience in this area.



### CUSTOM SOLUTIONS

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



STORAGE SOLUTIONS

**VSCALER®**



Boston have partnered with vScaler to provide hyper-converged infrastructure as a service using 'appliance building blocks' that dramatically simplify the deployment and support of your own private cloud. With no vendor lock-in and multiple service options to meet your budget and requirements, we deliver your cloud, your way. vScaler is based on OpenStack, and has

an intuitive interface for deploying Linux or Windows VMs and containers to bare-metal servers, configuring shared storage resources, managing your network, defining user security protocols, or monitoring system performance – giving you more time to focus on your users' applications.

**VSCALER**

**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 LINPACK 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 200 VMs in a single Hyper-Converged 2U appliance

**INTEGRATIONS FOR:**

-  Back-up
-  Recovery
-  Billing
-  Cloud Migration

**OPTIMISED FOR:**

-  GPUs
-  Low latency networking solutions
-  Scalable storage systems

**HYPERCONVERGED INFRASTRUCTURE FOR:**

-  Business
-  HPC
-  AI
-  Big Data

**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.

STORAGE SOLUTIONS

# DDN EXASCALER



For over 20 years, DDN has designed, developed, deployed and optimised systems, software and storage solutions that enable enterprises, service providers, universities and government agencies to generate more value and to accelerate time to insight from their data and information on-premise and in the cloud.

Developed and optimised using the latest advances in filesystem software technology, the DDN EXAScaler storage appliance delivers extreme performance, scalability, capability, reliability and simplicity.



## KEY FEATURES

### ENTERPRISE STORAGE FEATURES IN A HPC ARRAY

- Simplify data management and orchestration with API driven data integration tool.
- Snapshots, Encryption, Data Integrity, Quotas, Cloud Workload analytics (SIA), MultiCloud readiness & Multitenancy – all features that you would see in Enterprise grade storage.

### MULTI-PROTOCOL SUPPORT & FEATURES FOR EVOLVING CUSTOMER NEEDS

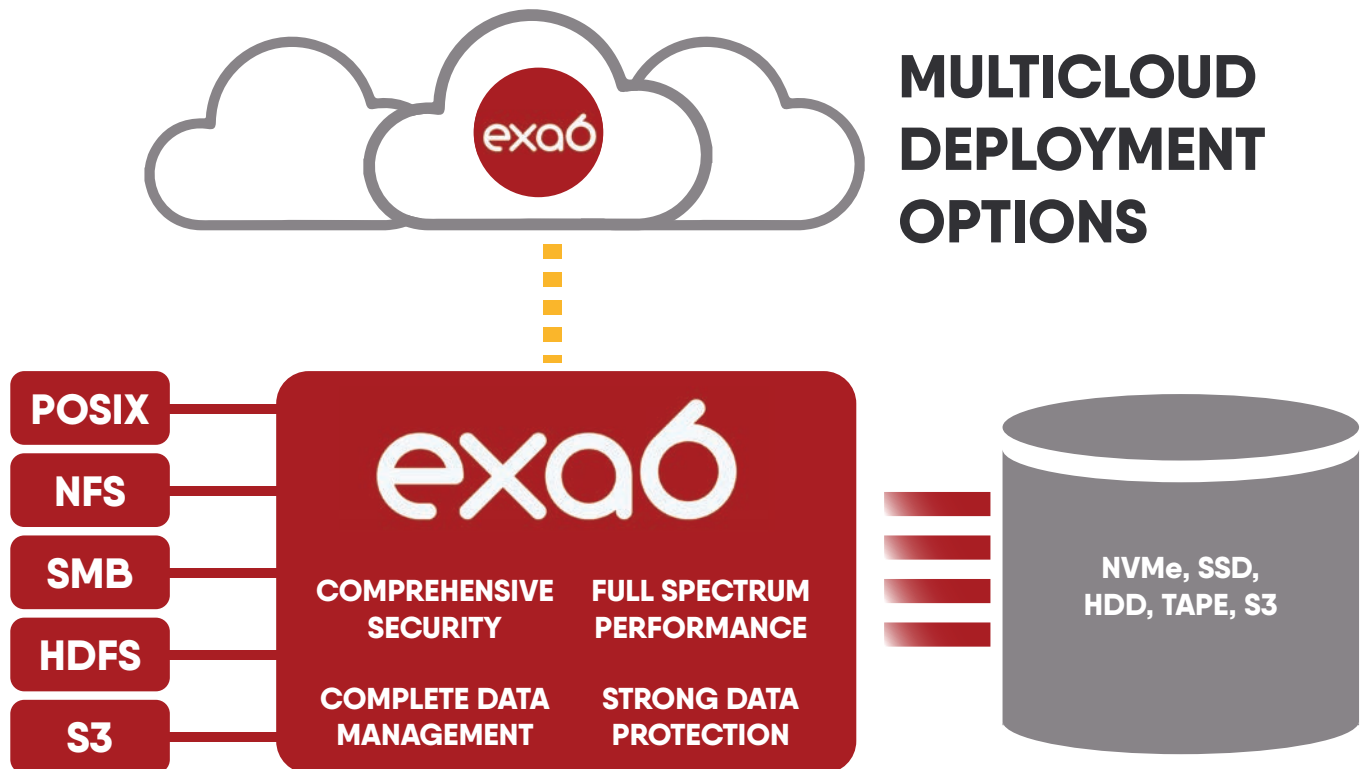
- File (NFS & SMB) Object (S3) and Container support
- Fastest throughput per SSD, scalable metadata all drive performance.

## DDN EXAScaler

EXAScaler introduces several new data management and integrity filesystem features developed by DDN and is only available in its appliances and cloud offerings. Stratagem is a powerful data orchestration engine that gives users comprehensive data residency controls using policy-based placement. Hot Pools intelligently moves data between high-performance flash and large capacity disk and ensures efficient use of storage. A native T10DIF implementation ensures that data is handled with full integrity from application to disk. Several dozen other new features in EXAScaler deliver unique value to users looking to deploy the most demanding workloads on premise, and in the cloud.



## MULTICLOUD DEPLOYMENT OPTIONS



## USE-CASES

### ARTIFICIAL INTELLIGENCE

Emerging AI use cases are driving growth-oriented data management strategies. Where traditional storage architectures fall short, EXAScaler is optimised to address the challenges of performance, scale and data mobility with simplicity and reliability.

### HIGH PERFORMANCE COMPUTING

Scientists and engineers are continuously pursuing new ways to solve complex problems faster. HPC workloads like physics modeling, oil and gas discovery & genomics process massive amounts of data accessed by multiple systems simultaneously.

## CASE STUDY

CHALLENGE	SOLUTION	ADVANTAGE
A leading life sciences university required shared computational resources to accelerate genomic medicine research, whilst keeping costs down under a tight budget.	System overhaul for improved performance and build capabilities for simulations, data analysis and sharing.	Enabled system overhaul under limited budgets, combining 200G HDR Infiniband and 40/10 Gbps Ethernet with a DDN Storage solution to improve performance and build capabilities, with up to 90 GB/sec performance per appliance.

STORAGE SOLUTIONS

# WEKA IO

Built for NVMe flash and cloud-native, Weka is a parallel file system designed to enable organisations to maximise the full value of their high-powered IT investments - compute, networking and storage. By leveraging existing technologies in new ways and augmenting them with engineering innovations, Weka's software delivers a more powerful and simpler solution

that would have traditionally required several disparate storage systems. The resulting software solution delivers high performance for all workloads (big and small files, reads and writes, random, sequential, and metadata heavy).

WEKA IO



## KEY FEATURES

### WEKA FS™

Discover a better way to harness the power of data to deliver better outcomes. Weka's Data platform was built from the ground up to shatter the limits imposed by legacy storage architectures by taking a unique software-only approach to address the needs of modern workloads. WekaFS is ideally suited for the challenges of mixed workloads—large and small files, random and sequential access, structured and unstructured data.

#### FIND THE RIGHT PLACE FOR YOUR DATA


Manage all your file and object data on NVMe, SSD flash or disk in the system. Leverage object connectivity to utilise low cost on-prem storage or the cloud.

#### EASY MANAGEMENT

Management is easy with a simple GUI. Rapidly provision, manage policies, change data protection levels and monitor system health intuitively.

#### ELIMINATE GUESSWORK

No tuning is required once set up. WekaFS is ideal for the challenges of mixed workloads. Files large and small, random and sequential, data structured and unstructured.



## USE-CASES

## ARTIFICIAL INTELLIGENCE

Artificial Intelligence requires fast data processing. Weka alleviates challenges across the entire data pipeline whether running across on-prem and the public cloud.

## CONTAINERS

Weka's solution for Kubernetes helps deliver value from the cloud to the edge with at scale performance.

## FINANCIAL SERVICES

Quantitative modelling helps mitigate risk, detect fraud and develop new trading algorithms. Weka covers latency sensitive workloads at exabyte scale.

## CASE STUDY

## CHALLENGE

Genomics England had previously implemented a scale-out NAS solution from a leading vendor to support the 100,000 genome project; however, it had already hit its limit on storage node scaling & performance suffered when the system was near capacity.

## SOLUTION

A two-tier architecture that takes commodity flash and disk-based technologies, presenting it as a single hybrid storage solution. The primary tier consists of 1.3 Petabytes of high performing NVME flash storage which supports the working data sets.

## ADVANTAGE

GE was able to realise no limit on capacity scaling, over 10x improvement in performance, 75% reduction in storage cost per genome, as well as now embedding a full disaster recovery strategy and offering integration with public cloud for compute elasticity.



STORAGE SOLUTIONS

# IBM SPECTRUM SCALE™



Enterprises and organisations are creating, analysing and keeping more data than ever before. In delivering those insights, an organisation's underlying information architecture must support the hybrid cloud, big data and artificial intelligence (AI) workloads along with traditional applications while ensuring security, reliability, data efficiency

and high performance. IBM Spectrum Scale™ is capable of managing petabytes of data and billions of files providing world class storage management with scalability, performance and policy-based storage tiering.



## KEY FEATURES

### SPECTRUM SCALE™

IBM Spectrum Scale™ is a parallel file system, where the intelligence is in the client and the client spreads the load across all storage nodes in a cluster, even for individual files. The IBM Spectrum Scale™ architecture allows it to seamlessly handle tens of thousands of clients, billions of files and yottabytes of data. IBM Spectrum Scale™ allows different applications or services to access the same data without moving or altering it. The system can span multiple storage environments and datacentres to eliminate data silos and "filer sprawl."

#### SIMPLIFIED DATA MANAGEMENT

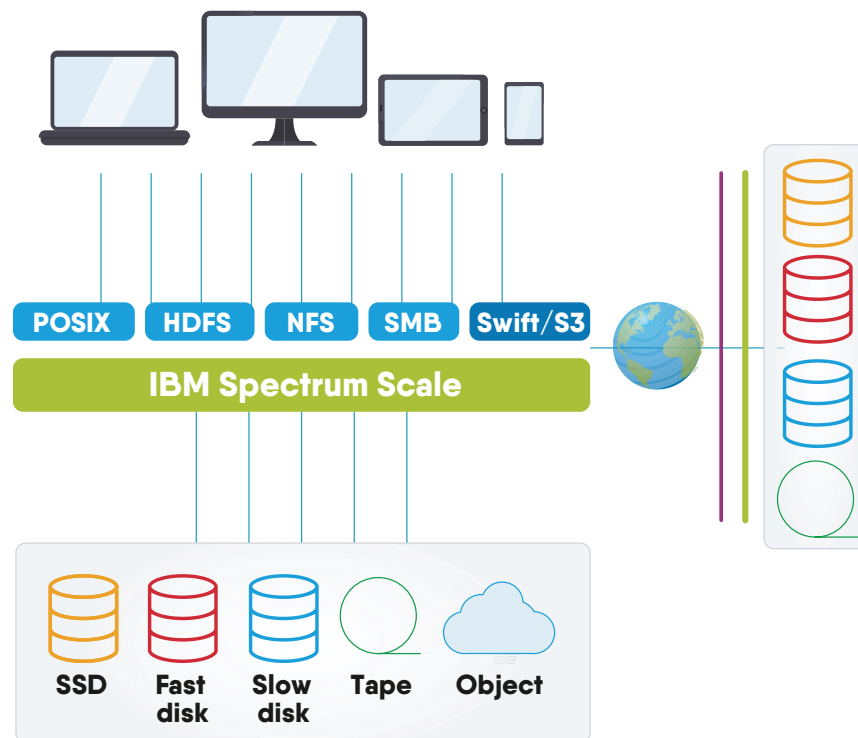
The Spectrum Scale client spreads load across all nodes in a cluster. This means that the architecture can scale easily while providing access to thousands of clients.

#### SIMPLE ADMINISTRATION

The enhanced GUI supports many features: performance, capacity, cloud tiering and enhanced maintenance with IBM remote support.

#### SECURITY ENHANCEMENTS

File audit logging capability tracks user access to the filesystem across all nodes and protocols. This ensures you have full visibility of who is accessing your storage.



## USE-CASES

### HIGH PERFORMANCE COMPUTING (HPC)

- The Spectrum Scale HDFS connector enables HPC customer to spin-up and terminate Hadoop or Spark clusters on their existing super computers like any other HPC job.

### DATA INTENSIVE APPLICATIONS

- Spectrum Scale enables scientists to seamlessly integrate HPC-like infrastructure into their experiments and into their workflows to get timely insight in new data sets.

## CASE STUDY

### CHALLENGE

To maintain its reputation as a premier research institution, a leading UK university needed to ensure that data is always available to a growing number of users running increasingly complex simulations.

### SOLUTION

The university deployed IBM® Spectrum Scale™ and IBM Spectrum Protect™, increasing transparency around data location and who accesses it, and increasing its mobility within a diverse IT environment.

### ADVANTAGE

The solution supplied supports compliance with data protection regulations at a low cost without disruption. Significant cost saving due to operational efficiency & 5000 researchers are supported by infrastructure that helps them get to results faster.

# COOLING SOLUTIONS



## LIQUID COOLING

Liquid cooling is rapidly becoming an inevitable approach to HPC cooling. Why is that? The short answer is that air cooling will not be able to keep pace with the wattage of components. The long answer is that on top of meeting temperature requirements, it also works towards sustainability goals as liquid cooling requires less energy to maintain, and generally there is less maintenance with a liquid cooled setup.

Using liquids to cool a system can come in many shapes and sizes. There are two categories which then also break into two types; single-phase and two-phase liquid cooling. Each approach can be applied directly on the chip or by immersing components, or even the entire system.

The single- or two-phase refers to the state of matter the cooling fluid is in. In a single-phase setup the liquid is heated up, stays a liquid, is cooled down elsewhere and looped back to start the process again.

A dual-phase setup makes the liquid evaporate on contact with the component, turning into a gaseous state. Once this vapour is cooled down, it condenses back into a liquid and starts the process again. Each have their advantages over the other, but they all excel in specific uses.





## COOLING SOLUTIONS FOR HPC




### ZUTACORE

The in-rack HyperCool™ solution uniquely supports up to 20kW computing power with an in-rack air cooled condenser and 100kW computing power with an in-rack water cooled condenser. This liquid cooling solution can enable customers to keep up with higher performance systems with shorter latencies, while saving scarce resources including energy, water and space. HyperCool™ is densifying datacentres from small to large scale, for both datacentre retrofits and new builds.



### ASPERITAS

With an impressive catalogue of expertise to ensure the solution meets your full requirements; the collaboration between award winning immersion cooling experts Asperitas and Boston offers an end-to-end immersion cooling “datacentre-in-a box” solution for enterprise, cloud and telecom providers, HPC and datacentre operators. Asperitas’ AIC24 immersion cooling solutions can facilitate optimised immersion ready servers with a max. 21” format.



### SUBMER

Submer helps enable datacentres that make sense through products, platforms, APIs, processes and installations that will move hyperscalers, colocation and huge industries to new levels of efficiency and innovation.

COOLING SOLUTIONS

# ZUTACORE

ZutaCore's innovative direct-on-chip, waterless, two-phase liquid cooling (2PLC) solution helps customers meet and surpass the challenges posed by high performance computing requirements. This solution goes beyond the limits of air and eliminates the use of water, mitigating the risk of IT failure and maximising cooling efficiencies. 2PLC solves thermal

challenges and provides significant datacentre power reduction. Furthermore, 2PLC is future proof and prepared for any evolution in high-powered chips: there is no limit to what it can cool as processors progress toward 1000W and even beyond.



## KEY FEATURES


### COOLED-BY-ZUTACORE

The Cooled-by-ZutaCore In-Rack solution fits into any rack in almost any environment. It uniquely supports up to 20kW computing power with an in-rack air cooled condenser and 100kW computing power with an in-rack water cooled condenser. The air-cooled HRU needs no special plumbing or external components outside of the rack (the water-cooled HRU does) – making it an ideal solution for existing data centers needing to accommodate higher power density in servers and racks as well as edge applications.

50%	3x	50%	0%	50%	>1000W
LESS ENERGY	PROCESSING CAPACITY	LESS SPACE	RISK OF MELTDOWN	CAPEX REDUCTION	CHIP COOLING

## COOLED BY ZUTACORE

The ZutaCore HyperCool waterless two-phase, liquid cooling (2PLC) based system provides a unique combination of benefits that directly addresses today's computing density needs and the future cooling demands of the world's computing infrastructure. Mechanically self-regulated, HyperCool provides on demand capabilities simply and reliably for high performance computing applications requiring high densities in new design and retrofits.



## TWO-PHASE, LIQUID COOLING (2PLC)

In-rack units combine the hardware system with a Software-Defined-Cooling (SDC) platform. By reducing the datacentre design to commissioning cycle and halving costs, ZutaCore empowers datacentre owners and operators to accelerate ROI and maximize real estate assets by supporting the proliferation of autonomous and central datacentres.

## BENEFITS

- Minimal additional space (RHx door or in-rack HRU).
- No chilled water-cooling system required to remove the high heat load.
- Easily deployed without modifying existing infrastructure and easy maintenance.
- Installation of high-powered computing nodes in an existing data center facility served by a traditional room cooling technology.
- The direct-on-chip system uses a non-conductive refrigerant, eliminating water inside the rack.
- Dissipates heat at the source to cut cooling power required from the server to the datacentre.
- Engineered for low-flow and low-pressure for light, compact design, and high density.

## WHY ZUTACORE?

### BIG IMPACT IN A SMALL FOOTPRINT

ZutaCore's Hypercool can be configured for use in racks, in rear doors and in whole rows.

### ENERGY EFFICIENT DATACENTRES

HyperCool slashes the consumption of scarce resources including energy, water and land.

### MORE COOLING IN LESS SPACE, WITH LESS ENERGY

HyperCool triples the processing capacity of highly dense computing environments.

### UNRIVALED HEAT DISSIPATION AT THE CHIP

HyperCool is uniquely built for dense, high temperature compute environments.



## COOLING SOLUTIONS

## ASPERITAS

Since 2014 Asperitas has worked on validating and developing Immersed Computing® as a unique approach to the datacentre industry. Building on existing liquid immersion cooling technologies by adding integration of power and network components, improving cooling physics with a strong focus on design and engineering for usability, Asperitas has come up with a complete and integrated solution which can be effectively utilised in most, if not all situations.

Asperitas is a high-tech company who also put emphasis on sustainability. As a leading immersion cooling specialist, they provide cutting edge solutions for energy efficient and high-density datacentres globally. Immersed Computing®,

Asperitas' unique award-winning natural convection-based technology concept, enables sustainable and high-density datacentres to run your applications anywhere they are needed.

Accelerate your infrastructure with application-specific hardware in fully integrated, contained and plug and play immersion cooling solutions for datacentres. The fact that Asperitas' solutions use convection to circulate the fluid makes it even more sustainable as there are no pumps involved, further lowering the overall total cost of operation.



## 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				

## COOLING SOLUTIONS

**SUBMER**

Submer's main goal is to make the operation and construction of data centres more efficient and sustainable. Submer uses their products, platform, APIs, processes, and installations to make hyperscalers and huge industries more efficient.

They put the needs of their customers first by providing solutions that are not only beneficial today but are future-proof also. Their solutions address questions like heat-reuse, net-zero water, and site selection. The single-phase immersion cooling Submer uses is flexible as it can be scaled without issue, all while staying space- and energy efficient.

Submer is in a unique position where they not only provide technology and solutions, but also develop them. They have expertise in thermodynamics, engineering and chemical engineering, to name a few.

With such capabilities, they are a forward-thinking company who designs solutions that address issues that may arise in the future as technology advances and our reliance on them increases.

**KEY FEATURES**

- Achieve a PUE as low as 1.03
- Up to 100kw of compute density
- Reduce latency and increase speed of deployment
- Increase hardware life-span
- Reduce hardware failure rate
- Save 50% on CAPEX building costs
- Minimize water usage

**BENEFITS**

Submer helps businesses optimise the operation of their data centre by leveraging their immersion cooling technology. Data centres that rely on air cooling must consider the flow of air and thereby manage hot and cold aisles. Not only is this additional effort from the employees, but it also takes up a lot of space. This approach to cooling has a slew of benefits which include space efficiency, higher potential density, and overall lower cost of operation. In turn, these benefits also trickle over to support sustainability goals, by lowering the impact on the environment.

# NETWORKING SOLUTIONS



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

## NETWORKING SOLUTIONS FOR HPC



### 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

NVIDIA Networking combine industry standard InfiniBand technology, supporting up to 400Gb/s in their adapter and switch products, with integrated InfiniBand Router and InfiniBand to Ethernet gateways, NVIDIA switches provide scalable fabric for powering the world's largest and fastest high-performance computing systems and next generation datacentres.



### ETHERNET SWITCHES AND ADAPTERS

NVIDIA Networking 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 NVIDIA Networking LinkX product family of cables and transceivers provides the industry's most complete line of 10, 25, 40, 50, 100, 200 and 400Gb/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 multiple form factors.



## NETWORKING SOLUTIONS

# NVIDIA INFINIBAND QUANTUM™-2 PLATFORM

The NVIDIA Quantum™-2 based QM9700 and QM9790 switch systems deliver an unprecedented 64 ports of NDR 400Gb/s InfiniBand per port in a 1U standard chassis design.

A single switch carries an aggregated bidirectional throughput of 51.2 terabits per second (Tb/s), with a landmark of more than 66.5 billion packets per second (BPPS) capacity.

Supporting the latest NDR technology, NVIDIA Quantum™-2 brings a highspeed, extremely low-latency and scalable solution that incorporates state-of-the-art technologies.



**NVIDIA**  
NETWORKING

## KEY FEATURES

### THE ERA OF DATA-DRIVEN COMPUTING

The QM9700 NDR InfiniBand switches extend NVIDIA In-Network Computing technologies and introduce the third generation of NVIDIA SHARP technology, SHARPV3.

### ENHANCED MANAGEMENT

The internally managed QM9700 switch features an on-board subnet manager that enables simple, out-of-the-box bringup for up to 2,000 nodes.

## PERFORMANCE

As high-performance computing (HPC) and artificial intelligence (AI) applications become more complex, the demand for the most advanced high-speed networking is critical for extreme-scale systems.

NVIDIA Quantum™-2 is the industry-leading switch platform in power and density, with NDR 400 gigabit per second (Gb/s) InfiniBand throughput that provides AI developers and scientific researchers with the highest networking performance available to take on the world's most challenging problems.

NETWORKING SOLUTIONS

# NVIDIA ETHERNET SPECTRUM®-4 PLATFORM

The world's first 400Gbps end-to-end networking platform, NVIDIA® Spectrum®-4 provides 4x higher switching throughput than previous generations, with 51.2 terabits per second.

It consists of the NVIDIA Spectrum®-4 switch family, NVIDIA ConnectX®-7 SmartNIC, NVIDIA BlueField®-3 DPU and the DOCA™ data center infrastructure software to supercharge cloud-native applications at scale.

Built for AI, NVIDIA® Spectrum®-4 Switch arrives as datacentres are growing exponentially and demanding extreme performance.



## KEY FEATURES

### FLEXIBILITY AND CONTROL

Deploy in your layer-2 and layer-3 cloud designs, in overlay-based virtualized networks, or as part of high-performance, mission-critical ethernet storage fabrics. In any deployment, you'll get: Industry-leading performance and predictability, comprehensive visibility, high-performance remote direct-memory access over converged ethernet (RoCE), and advanced virtualization.

### END-TO-END OPEN NETWORKING

Open networking is about offering choice and flexibility while removing restrictions, and that means being open at all tiers of the stack. NVIDIA is the leader in end-to-end open networking at all layers of software and hardware.

## PERFORMANCE

NVIDIA Spectrum®-4 features a fully-shared and monolithic packet buffer that's dynamically available to all ports. This provides excellent microburst absorption with true, port-to-port, cut-through latency. Spectrum®-4 also supports programmability of the pipeline and packet modifier/parser without impact to latency or packet rate, enabling flexibility without compromising performance.

The on-chip packet buffer supports high-bandwidth packet reads/writes, enabling a consistent and predictable performance. The monolithic packet buffer architecture simplifies buffer management and traffic scheduling, while also enabling fair resource sharing.



# SERVICES



## BOSTON TRAINING ACADEMY

The mission of the Boston Training Academy (BTA) is to become a renowned developmental ground for talent engagement, education and solutions across a variety of disciplines.



## BOSTON TECHNOLOGY CONSULTING

Our technical and HPC experts, based around the world, are constantly benchmarking new technologies, 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.



## 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 SERVICES & SUPPORT

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





## EXPERT AI AND DEEP LEARNING TRAINING FOR YOUR STAFF AND CUSTOMERS



Boston Training Academy (BTA) provides structured training by world-class trainers who deliver tailored courses based on the current knowledge of the attendees. Our program content is designed by industry experts and covers data science skills, exposure to analytical tools, corporate learning, consulting on AI projects, to name a few.

The Centre of Excellence (CoE) at BTA brings together experts from industry academia and Government to provide the latest infrastructure, smart tools, skilled resources and leadership for AI exploration. The CoE-AI aims to create the most sophisticated ecosystem for AI, Data Sciences, Machine Learning and Deep Learning to enable breakthrough innovations.



BTA also offers online starter courses for complete beginners to get working knowledge on the fundamentals of NLP, computer vision, deep learning, and Python for data science. The Deep Learning Institute (DLI) is a workshop by NVIDIA, hosted by Boston where participants get hands-on training sessions, with help from developers, data scientists, and researchers. Attendees will learn how to approach challenges using deep learning techniques like building transformer-based NLP, building conversational AI applications, building intelligent recommender systems, among other approaches.

# 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





## EFFECTIVELY CREATE, OPTIMISE AND GROW YOUR BUSINESSES WITH EXPERT CONSULTING



Boston Technology Consulting (BTC) helps organisations to effectively create, optimise, and grow their business in today's data and IT focused market. Our consultants mastermind solutions for the most complex business needs in an effort to achieve their goals, while balancing costs and performance. Our core values are leadership, integrity, collaboration, accountability, and simplicity.



Our consultancy services cover a variety of fields and verticals: SaaS digital applications, cloud and IT infrastructure, data and analytics, cyber security, audit and compliance, digital transformation, Emerging technologies (EmTech).

BTC with its global reach equips us with the unique skill to provide agile and scalable solutions to everyone from the smallest players to industry heavy hitters. BTC's main objective is to solve current problems while minimising risks of future issues. Such all-encompassing approach to problem solving has been utilised in the banking, financial services and insurance (BFSI), government, manufacturing, oil and gas, logistics, telecoms, pharma and retail.



# WORLD-CLASS BOSTON LABS

Boston Limited has been a staple of HPC hardware and software solutions for the past 30 years. Dealing with cutting-edge, state-of-the-art equipment has been an everyday occurrence, furthermore, being vendor-agnostic, we can provide you with solutions tailor-made for your needs and budget.

01

## LEVERAGE OUR EXPERTISE

Boston Labs is at the forefront of testing each and every component under stress to get an insight on how they perform in real-life environments. The components are also compared with previous models and with other vendors' products to give an unbiased review of a parts performance.

02

## EXPERT HELP EVERY STEP OF THE WAY

03

## AT YOUR CONVENIENCE WITH ANYTIME REMOTE DIAL-IN ACCESS

At Boston Labs you are given the unique opportunity to have a hands-on experience with components or entire systems guided by one of our experienced field application engineers. Depending on the solution you need, it might be arranged that you are given access to test the software that will upgrade your business so you would not even need to come in person.

04

## BE THE 1ST TO ACCESS THE LATEST AND FUTURE TECHNOLOGIES

Our line of hardware and software available for testing is always being updated, so whenever you want to test something, you know it is the best of the best and we are also ensuring that by fully stress-testing components. At Boston Labs every single component goes through rigorous testing. When new components are released, Boston Labs is some of the earliest testers and their findings can be read in their blog posts.

05

## 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

Boston Technology Consulting brings our expert consultants to address client's complex business needs and achieve their organisational and business goals. Our business partners and clients believe in our unique value proposition through use of cutting-edge technology and cost-effective solutions.

## 02 DESIGN

## 03 ASSEMBLE

## 04 CONFIGURE

## 05 TEST

## 06 SHIP

## 07 INSTALL

## 08 SUPPORT

With over 30 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 [www.boston.co.uk/about/services/default.aspx](http://www.boston.co.uk/about/services/default.aspx)

# ADDITIONAL SERVICES AND 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.

## ON-SITE WARRANTY PACKAGES

	RETURN TO BASE (RMA)	ADVANCED WARRANTY	ONSITE WARRANTY			
			PEWTER	BRONZE	SILVER	GOLD
TECHNICAL SUPPORT DESK	YES	YES	YES	YES	YES	YES
ADVANCED PARTS SHIPPING	NO	YES	NO	YES	YES	YES
REMOTE ISSUE TRIAGE	YES	YES	YES	YES	YES	YES
ONSITE SUPPORT/DIAGNOSIS *	NO	NO	NO	NO	NO	NO
ONSITE BREAK/FIX LABOUR	NO	NO	YES	YES	YES	YES
ENGINEER TO SITE TARGET (POST DIAGNOSIS)	N/A	N/A	NBD	NBD	SBD	4 HOURS
ACCESS TO SERVICES 24 HOURS A DAY & WEEKEND AND UK BANK HOLIDAYS	NO	NO	NO	NO	NO	YES




# TECH TALK

## PROFITS THROUGH SIMPLICITY

### THE CHALLENGE

A global cloud-based AI company wanted to migrate from public cloud to their own private cloud, to lower costs by having more access to management of hardware and security. The downside of public cloud like AWS or Google is that the user has limited to no access to the hardware, extra costs from renting, and less flexible data management.



### THE SOLUTION

Boston Limited provided an end-to-end service where the hardware from a mix of our vendors, the shipping, and the installation was all carried out by Boston. The main infrastructure is built on NVIDIA DGX which has versatility in AI training, inferencing, and analytics – this way the company did not have to commit to one aspect of AI workloads.

Supermicro and AMD CPU servers were also set up to accelerate the workload carried out by the NVIDIA DGX. The storage itself is Boston's very own Ceph-based, massively scalable, open source, software-defined storage system. All these components are connected via NVIDIA Mellanox switches.



### THE OUTCOME

After installation, the company saw a ROI in 6 months. This is thanks to a few factors: renting cloud storage is expensive, so by owning their own hardware, they already reduced their expenses just to maintain their business. Secondly, the company, and thereby the customers, enjoy the benefits of tighter security, as the security protocols are not dictated by the cloud storage provider. Finally, they have seen an increase of workload performance by 30%, this also translates to customer satisfaction as their own workloads became 30% more efficient.



# TECH TALK

## COMPUTE EXPRESS LINK™ (CXL™)

### WHAT IS IT?

Compute Express Link™ (CXL™) is a breakthrough high speed CPU-to-Device interconnect which has become the new industry I/O standard. CXL™ builds upon and supports PCI Express 5.0, utilising both its physical and electrical interface which allows for disaggregated resource sharing to improve performance whilst lowering costs.

CPU core counts have rapidly increased over the last decade, and this is set to continue with everyone now adopting chiplet architectures. However, the memory channel bandwidth per core has not been keeping pace and is falling behind. This is despite efforts to add channels by use of larger sockets, thicker PCB layer boards and bigger form factors. We have now come to physical constraints on all these fronts.

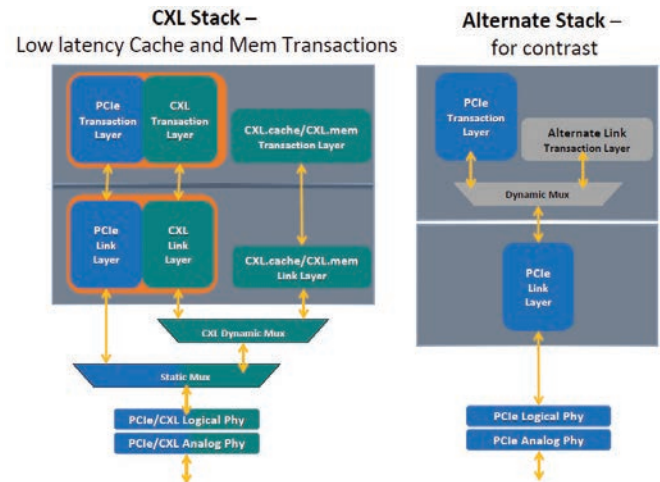
### HOW DOES IT WORK?

Compute Express Link™ is an open industry standard that leverages the existing and ubiquitous PCIe infrastructure.

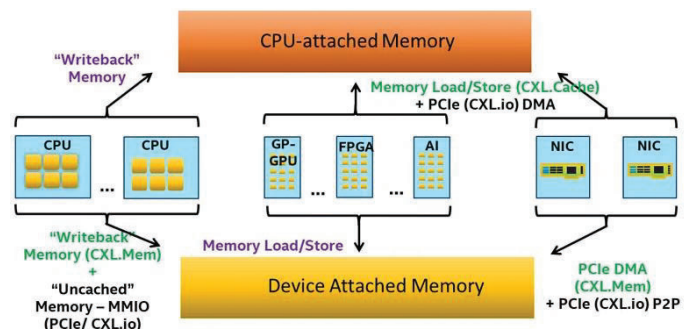
This enables a high bandwidth and low latency interconnect between CPUs and devices, it will also allow memory coherency between CPU's memory and attached devices' memory. This therefore allows resource sharing for optimum performance and lowering overall system costs.

### DYNAMICALLY MULTIPLEXED SUB PROTOCOLS

The Compute Express Link™ transaction layer is comprised of three dynamically multiplexed sub-protocols on a single link; CXL.cache, CXL.memory and CXL.io. CXL.io is used for link



link initialisation and management, device discovery, enumeration, interrupts, DMA and register access. CXL.memory protocol is responsible for providing a Host processor with direct access to device attached memory using load/store commands. CXL.cache protocol on the other hand allows for a CXL™ device – like a CXL™ compatible accelerator to access and cache processor memory.





**UK (Headquarters)**

+44 (0) 1727 876 100  
sales@boston.co.uk  
www.boston.co.uk

**FRANCE**

+33 (0)1 72 81 06 76  
sales@boston-it.fr  
www.boston-it.fr

**INDIA**

+91 22 5002 3262 (Mumbai office)  
+91 80 4308 4000 (Bangalore office)  
sales@bostonindia.in  
www.bostonindia.in

**AFRICA**

+27 826 486 485  
sales@boston.co.uk  
www.boston.co.uk

**GERMANY**

+49 (0) 89 9090 1993  
sales@boston-it.de  
www.boston-it.de

**SWITZERLAND**

+47 (0) 71 554 22 75  
sales@boston-it.ch  
www.boston-it.ch

**Powerful Technology**  
*Intelligently Applied™*