



HIGH PERFORMANCE COMPUTING



CONTENTS

WHAT'S INSIDE?

05	Software
06	Nvidia Base Command Manager
08	Warewulf
09	XCAT
s 10	Boston Services
11	Boston Technology Consulting
12	Boston Training Academy
13	Boston Labs
14	Professional Services
15	Support and onsite warranty
16	
17	
18	
19	
20	
22	
24	
26	
28	
28	
33	
34	
35	
36	
37	
38	
39	
40	
	05 06 08 09 10 11 12 13 14 15 16 17 18 19 20 22 24 26 28 33 34 35 36 37 38 39 40





INNOVATIVE SOLUTIONS

We innovate tech solutions for diverse sectors, adapting to evolving business needs



WHAT SETS US APART?

BESPOKE CUSTOMISATION

Tailored, bespoke solutions harmonise with customer objectives, defining our excellence in technological delivery



GLOBAL REACH

Our global presence assures swift support, regardless of location, backed by a history of successful collaborations





A proven track record with diverse projects instils confidence in us as a dependable partner



MANUFACTURING AND TESTING

Our in-house design, manufacturing and testing in the Boston Labs guarantees quality and quick turnaround



RESEARCH AND DEVELOPMENT

Ongoing R&D equips us with expertise for the latest tech advancements, providing customers a competitive edge





We are environmentally conscious and prioritise eco-friendly, sustainable solutions



ONE-STOP SHOP

Our diverse tech offerings, from computing to AI, position us as a one-stop shop for all your needs



CUSTOMER-CENTRIC

Customer-centricity is our cornerstone, understanding challenges and exceeding expectations in the solutions we create



BOSTON LIMITED YOUR TRUSTED PARTNER

WHO ARE WE?

BOSTON: THE POWER OF HIGH-PERFORMANCE

Boston Limited is a leading provider of high-performance, power-optimised technologies for various markets. We offer bespoke award-winning solutions that are designed to help our customers stay ahead of the curve.

Our products include servers, storage, workstations, and clustered solutions. We also provide a range of value-added services, such as global on-site maintenance packages and life cycle management for IT equipment.

What sets Boston apart is our expertise in integrating, testing, and validating highperformance custom solutions. We can fully brand and package our solutions to meet the specific needs of our customers.

We also have extensive experience in AI and data. We offer end-to-end services, from training and advisory services to implementation, to help organisations navigate the complexities of digitalisation successfully.

Our expertise in AI and ML allows us to provide solutions that automate processes, optimise operations, and provide valuable insights from data.

Our broad range of capabilities, commitment to customised solutions, and valueadded services have made us a trusted partner for organisations in various industries and markets worldwide. We recently celebrated our 30th anniversary, and we are expanding to meet the growing demand for high-performance systems.

If you are looking for a high-performance technology partner that can help you stay ahead of the curve, then Boston is the right choice for you.

 (\bigcirc)

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.







COMPUTE SOLUTIONS



SUPERMICRO

Supermicro's compute solutions offer high-performance server systems, including rackmount and blade servers, for scalability, reliability and energy efficiency in datacentres and cloud computing.

(0)



AMD

AMD's compute solutions feature powerful Ryzen processors for gaming and content creation, as well as EPYC processors for datacentres and enterprise workloads.



AMPERE

Ampere empowers enterprise computing with their high-performance AmpereOne processors, offering exceptional core counts, clock speeds and memory capacity.



INTEL

Intel's compute solutions include Core processors for a balance of performance and power efficiency, and Xeon processors for enterprise computing with scalability and advanced security.



NVIDIA

NVIDIA's Grace Hopper Superchip redefines large-scale high-performance computing and AI deployments, offering unrivalled computational capabilities.





SUPERMICRO COMPUTE SOLUTIONS FOR HPC

Supermicro and Boston have 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 customers' specific needs.

Reference designs already exist from Supermicro for enterprise HPC, scientific research, AI and deep learning – all of which can be tailored to the finest 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 datacentres around the world use Supermicro systems due to their reliability, density and ease of application. Additionally, 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 and 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 its 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.







AMD 4[™] 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 4th generation EPYC[™] 9004 series CPU.

The EPYC[™] family of CPUs provides great performance and security while maintaining power and cost-efficiency. AMD's 4th 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%.

KEY FEATURES:

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



- 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 4th generation EPYC[™] 9004 series processors carries on this tradition







AMPERE



Ampere is at the forefront of enterprise computing, redefining performance with their innovative AmpereOne processors. These state-of-the-art processors have a single-socket configuration with the potential for up to 192 cores, delivering exceptional clock speeds of up to 3.0 GHz and supporting up to 250W or 400W Thermal Design Power (TDP), depending on the model. The AmpereOne family processors provide an unprecedented 16 DIMM slots for memory capacity, ensuring your workloads run seamlessly and efficiently.

Whether you're looking for an enterprise server, high-performance computing or a Telco Edge solution, Ampere's AmpereOne processors provide the performance, power and memory capacity required to meet your demands. AmpereOne processors guarantee outstanding performance, allowing customers to fully explore and utilise their computing environments.

KEY FEATURES:

- Up to 192 cores, providing exceptional processing power
- Clock speeds of up to 3.0 GHz for swift and responsive computing
- Choose between 250W or 400W TDP for tailored performance
- Support for 16 DIMM slots ensures efficient memory management

- Ideal for enterprise, highperformance computing and Telco Edge solutions
- Tailor the power consumption to your specific computing needs
- Robust memory capacity ensures seamless workload handling





intel.

INTEL® 4TH GEN 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).

KEY FEATURES:

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



BENEFITS:

6

- New standards in datacentre architecture
 - Multi tile SoC for scalability
 - Physically tiled, logically monolithic
 - General purpose and dedicated acceleration engines
- Designed for Cloud, Microservices & Al workloads
 - Performance core architecture

RON

- Workload specialised acceleration



NVIDIA GRACE HOPPER SUPERCHIP

Grace Hopper stands as NVIDIA's foundational building block for expansive HPC and AI deployments, fusing CPU, GPU and memory technologies. What sets it apart is the exceptional accessibility up to 150TB of peer memory, made possible by combining 256 Superchips. This peer memory design eliminates the bottlenecks typically associated with traditional Ethernet networking and InfiniBand-based distributed memory used in current deployments. Grace Hopper redefines the landscape of data processing, ensuring efficient and seamless large-scale computing and AI operations.

With remarkable memory capacity and bandwidth, the Grace Hopper Superchip propels data centre efficiency and agility to unprecedented levels, enabling faster insights and real-time decision-making. This innovative solution represents a pivotal step forward in the world of computation, setting new standards for performance and adaptability across a multitude of industries.

KEY FEATURES:

- Extensive memory capacity and bandwidth, providing swift access to large datasets
- Supports seamless scalability for diverse workloads, from AI to scientific computing
- Offers a unified solution for AI and data-intensive tasks
- Streamlines data centre operation`s

- Elevates overall computing performance
- Brings together the power of CPU and GPU for unified AI capabilities, reducing bottlenecks in AI workflows
- Offers adaptability to various workloads and scales
- Experience heightened efficiency, driving faster insights and enabling real-time decision-making





ACCELERATED SOLUTIONS

AMD

AMD offers MI series GPUs for exascale compute in HPC, AI and machine learning workloads.

GRAPHCORE

GRAPHCORE

Graphcore developed the IPU, a specialised processor for AI compute, enabling new advancements in machine intelligence research.



INTEL

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.



NVIDIA

NVIDIA SuperPOD is a cutting-edge, powerful computing infrastructure that harnesses the immense processing capabilities of NVIDIA GPUs to accelerate artificial intelligence and high-performance computing workloads.



AMD INSTINCT[™] MI250X SERIES PLATFORM



AMD's new Instinct MI250X is a powerful datacentre GPU designed for server and supercomputer workloads, including some of the largest exascale systems. With CDNA2 architecture and 6nm lithography, the MI250X features 14,080 stream processors and 220 compute units, offering up to 383 TFLOPs of mixed precision FP16 performance for deep learning training. The GPU also delivers up to 47.9 TFLOPs of double precision (FP64) performance and 95.7 TFLOPs of single precision matrix (FP32) performance.

The MI250X supports a range of technologies, including AMD Infinity Fabric and ROCm software ecosystem, enabling developers to accelerate their time to science and discovery. It is designed for server platforms and is compatible with Linux x86_64. The GPU includes 128 GB of HBM2e memory with full-chip memory ECC support and up to 3276.8 GB/s peak memory bandwidth.

KEY FEATURES:

- CDNA2 architecture and 6nm lithography
- 14,080 stream processors and 220 compute units
- Up to 383 TFLOPs and mixed precision FP16 performance
- Support for AMD Infinity Fabric and ROCm software ecosystem

- Designed for server and supercomputer workloads, including exascale systems
- High-performance deep learning training with up to 47.9 TFLOPs of FP64 performance
- Accelerated time to science and discovery for developers
- Ample memory with 128 GB of HBMe and peak memory bandwidth of up to 3276.8 GB/s



GRAPHCORE

GRAPHCORE



Graphcore 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.

KEY FEATURES:

- 1.4 petaFLOPS of AI performance
- Slim 1U server blade with extreme Al power
- Supercomputing scale

- Increases power efficiency and performance
- Reduces training and inferencing times
- Al and machine learning capabilities







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:

- High-efficiency deep learning capabilities
- Increased object detection in autonomous vehicles, medical imaging and manufacturing

- 40% better price performance in the AWC cloud with Amazon EC2 DL1
- Lowers training times to make the previous generation more cost-efficient
- Boost performance through efficiency





NVIDIA® SUPERPOD



The DGX SuperPOD is an optimised data centre rack configuration containing five or more DGX[™] servers, plus networking switches to support single and multi-node AI model training and inference using NVIDIA AI software.

The DGX SuperPOD is also designed to be compatible with leading storage and networking technology providers. Boston offers a portfolio of NVIDIA DGX SuperPOD reference architecture solutions with leading storage vendors including: DDN, IBM Spectrum Storage, VAST Data, and WEKA. All incorporate the best practice of NVIDIA DGX SuperPOD design and are delivered as fully-integrated and ready-to-deploy solutions to make your data centre AI deployments simpler and faster.

KEY FEATURES:

- A compute fabric that is rail optimised in a full fat-tree topology, maximising capability for the DGX[™]
- Scalable, with the ability to combine multiple DGX SuperPODs and scale to thousands of nodes
- A storage fabric that maximises bandwidth providing more than 40GB/s for each DGX node. Storage fabric leverages RDMA for the fastest, low latency performance across the DGX SuperPOD and throughout the data centre

- Tested and proven end-end hardware and software platform
- Advanced infrastructure software that always improves
- Dedicated expertise and services
- Unleash productivity with NVIDIA base command manager



STORAGE SOLUTIONS



BOSTON IGLOO AI+

The Boston Igloo AI+ is Storage Made For AI. It has been designed specifically for the high performance needs and low budget requirements of AI projects - without compromise.



DDN EXASCALER

DDN Exascaler offers appliances delivering fast, reliable and efficient data storage and retrieval, which makes them ideal for businesses that require high-speed data processing.



WEKA

IBM

IBM offers storage solutions, including flash storage, hybrid cloud storage and data protection solutions. Their all-flash arrays deliver high-speed data transfer and low latency, making them ideal for businesses that require real-time data access.

WEKA

Weka provides storage systems that are designed for data-intensive applications and workloads, providing high-speed data access and processing that can keep up with demands.





BOSTON IGLOO AI+

Investment in accelerator technology can be critical for the success of an AI project. We have designed the Igloo AI+ with blistering performance to support AI workloads without the need to compromise investment into accelerators. The NFS protocol allows multiple DGX servers to access the storage meaning that more systems can be connected to the environment as the need for AI compute grows.

The Boston Igloo AI+ strikes the balance of cost versus blistering performance making it the perfect fit for storage needs that generally run into the low hundreds of TB storage capacity needed to feed the accelerators with data to develop a model in as short a time as possible. For smaller environments with a single GPU system, the Boston Igloo AI+ can be connected directly WITHOUT the need for a switch! This furthers investment that can be made into GPU's!

KEY FEATURES:

- Block or file access
- Designed for AI workloads
- Up to 80GB/s+ throughput
- RAID 1, 10, 5 and 6 support
- GPU direct support
- Up to 737TB capacity

- Exceptional value for money
- Packaged appliance Built, configured and remote installation service from Boston
- NFS protocol allows additional DGX systems to be added over time
- Set it and forget it. No need for intensive administration



SPECIFICATIONS

MODEL	IGLOO AI+ 1U FORM FACTOR	IGLOO AI+ 2U FORM FACTOR		
Drives	6 or 12x NVMe	6, 12 or 24x NVMe		
CPU	(2x) Intel Gold 6442Y	(2x) Intel Gold 6442Y		
Memory	256GB	256GB		
Network Adapter	(2x) NVIDIA 200GbE card	(2x) NVIDIA 400GbE card		
Installation	Remote Installation	Remote Installation		
Capacity	Up to 368.6TB RAW	Up to 737.3TB RAW		
Throughput	40GB/s+	80GB/s+		

USE CASE:

CHALLENGE

Many organisations face a choice when buying storage to support NVidia DGX technology. They wish to invest as much as possible into DGX only to find that little of the budget remains for fast parallel storage. Enterprise All Flash does not deliver enough performance, while a traditional parallel storage system is too expensive.

SOLUTION

The Boston Igloo AI+ meets the need by providing blistering storage performance at a pricepoint that can support the limitations of a small budget. The NFS filesystem means that DGX systems can be added to storage over time while not comprimising on performance needed to support the AI workloads.

ADVANTAGE

The Boston Igloo AI+ has been designed especially for this common use case. The low cost of this fast storage solution means that on many occasions budget can be directed into more GPUs.



DDN EXASCALER



For decades, 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.

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.

KEY FEATURES:

- Enterprise storage features in 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

- Multi-protocool support and features for evolving customer needs
- File (NFS & SMB) Object (S3) and Container support
- Fastest throughput per SSD, scalable metadata all drive performance





USE CASE:

CHALLENGE

A leading life sciences university required shared computational resources to accelerate genomic medicine research, while keeping costs down under a tight budget.

SOLUTION

System overhaul for improved performance and build capabilities for simulations, data analysis and sharing.

ADVANTAGE

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.

23



IBM® SPECTRUM STORAGE™

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 Storage[™] is capable of managing petabytes of data and billions of files providing world class storage management with scalability, performance and policy-based storage tiering.

IBM[®] Spectrum Storage[™] 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 Storage[™] architecture allows it to seamlessly handle tens of thousands of clients, billions of files and yottabytes of data.

KEY FEATURES:

- The enhanced GUI supports many features: performance, capacity, cloud tiering and enhanced maintenance with IBM remote support
- File audit logging capability tracks user access to the filesystem across all nodes and protocols
- The Spectrum Storage[™] client spreads load across all nodes in a cluster. This means that the architecture can scale easily while providing access to thousands of clients

- Simplified data management
- Simple administration
- Security enhancements



USE CASE:

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 Storage[™] 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 savings due to operational efficiency, and 5000 researchers are supported by infrastructure that helps them get to results faster.



WEKA IO



WEKA

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'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.

KEY FEATURES:

- Manage all your file and object data on NVMe, SSD flash or disk in the system within a single global namespace
- Management is easy with a simple GUI. Rapidly provision, manage policies, change data protection levels and monitor system health intuitively
- No tuning is required once set up. WekaFS is ideal for the challenges of mixed workloads

- Manage all your data from one, centralised location
- Eliminate guesswork
- User friendly management



USE CASE:

CHALLENGE



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.





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.

DIFFERENT TYPES OF LIQUID COOLING

DIRECT LIQUID COOLING (DLC)

In the realm of liquid cooling, single-phase DLC maintains a consistent liquid state, efficiently absorbing heat from components and subsequently dissipating it through cooling mechanisms before being recirculated, finding applications in both direct-on-chip and immersion cooling setups.

TWO-PHASE DIRECT LIQUID COOLING (2-PHASE DLC)

Two-phase DLC involves an intricate process where a liquid undergoes a phase change into vapour upon contact with heat and, after vapour cooling, it condenses back to its liquid form, a method often employed in immersion cooling; it relies on a precise choice of a fluorocarbon-based liquid that transitions states around 50°C, demanding careful handling.

IMMERSION COOLING

Immersion cooling submerges components within a dielectric fluid, offering the choice of both single and two-phase cooling loops, its dielectric properties ensuring safety as it efficiently absorbs and dissipates heat from electronic parts; with a secondary heat exchanger in play, the captured heat can be redirected for dissipation or repurposing, making it particularly well-suited for high-power applications such as datacentres and industrial processes.



DIRECT LIQUID COOLING













SINGLE-PHASE IMMERSION





LIQUID COOLING



ASPERITAS - SINGLE PHASE IMMERSION

Asperitas' solutions use convection to circulate the fluid making them even more sustainable as there are no pumps involved, further lowering the overall total cost of operation.



SUBMER - SINGLE PHASE IMMERSION

Submer 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.



SUPERMICRO - DLC SINGLE PHASE

Supermicro offer a range of liquid cooling integrated solutions. Direct to chip liquid cooling solutions deliver superior performance and reduce datacenter electricity costs by up to 51%.



ZUTACORE - DLC 2-PHASE

Zutacore's solutions enable organisations to elevate their providings through compute-and data-intensive applications, ensuring their systems' best performance and longevity.







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.

KEY FEATURES:

- Reduces CO2 emissions by up to 40%
- 98% energy reuse
- Reduces TCO by up to 40%

- Plug and play capabilities
- Modular system
- Increases density by up to 5x



SUBMER



Submer's main goal is to make the operation and construction of datacentres 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.

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
- Minimise water usage





SUPERMICRO

The Supermicro liquid cooled rack solution delivers superior performance and efficiency for HPC, large scale AI, and cloud scale compute infrastructure.

Servers that are application optimised for AI, HPC, and Analytics require the latest in CPU and GPU technologies, which run hotter than previous generations. Multiple CPUs and GPUs per server are needed for performance intensive computing, driving up the REDUCTION electricity demands for the server as well as at the rack level.

Supermicro's liquid cooled rack solution is made of several components that are designed inhouse to achieve the highest level of performance and reliability. All the components are integrated as a rack level solution providing a true one-stop shop customer experience.

KEY FEATURES:

- Enterprise grade redundant cooling pump and power supplies, leak-proof connectors, and leak detection
- Options for Processor and GPU cooling

- Full turn-key single source solution from Supermicro
- Accelerated lead times based on in-stock inventory
- Up to 51% reduction in electricity costs in the datacenter





ZUTACORE



Having optimal server capabilities are essential for organisations engaged in compute- and data-intensive applications. These organisations must invest in efficient cooling solutions to ensure their systems' best performance and longevity.

This investment is especially true as each successive generation of processors creates more heat and has more stringent cooling tolerances. The ZutaCore dielectric liquid cooling solution is one of the most effective methods for providing reliable cooling to high-end processors. It has been proven to enable CPUs to perform consistently at high utilisation rates with zero throttling.

KEY FEATURES:

- 50% less energy
- 3x Processing Capacity
- 50% less space

- >1000W of Chip Cooling
- Waterless dielectric liquid cooling
- 50% Capex Reduction





NVIDIA NETWORKING



NVIDIA ETHERNET

NVIDIA Networking Ethernet products provide scalable, low-latency and efficient fabric with simplified management, reducing overall costs and power consumption compared to traditional Ethernet.



NVIDIA INFINBAND

NVIDIA Networking's InfiniBand switches and adapters support up to 400Gb/s, enabling scalable fabric for high-performance computing and datacentres.



NVIDIA SPECTRUM X

The NVIDIA Spectrum[™]-X networking platform is the first Ethernet platform designed specifically to improve the performance and efficiency of Ethernet-based AI clouds.



 \bigcirc



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[™] datacentre 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.

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.

KEY FEATURES:

- 400 Gps end-to-end networking platform
- 4x higher switching throughput

BENEFITS:

- Provides extreme performance
- Enhanced efficiency
- 4x higher switching throughput than previous generations

• 51.2 TB/s





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.

KEY FEATURES:

- High speed, low latency and scalable
- Aggregated bidirectional throughput of 51.2 TB/s
- Delivers 64 ports of NDR 400 GB/s

- Incorporates state-of-the-art technologies
- 66.5 BPPS capacity
- 1U standard chassis design





NVIDIA SPECTRUM X

Spectrum-X is a revolutionary solution created by NVIDIA for building multi-tenant, hyperscale AI clouds and AI Factories with Ethernet. With it, organisations can significantly improve the performance and power efficiency of AI clouds and gain higher predictability and consistency, resulting in faster time to market and a greater competitive edge.

NVIDIA Spectrum-X enhances performance and power efficiency by up to 1.7X, accelerating the processing, analysis, and execution of AI workloads and, in turn, the development and deployment of AI solutions. Spectrum-X is fully standards-based Ethernet with support for open Ethernet stacks (SONiC) at cloud scale, and is tuned and validated across the full stack of NVIDIA hardware and software, creating an unmatched Ethernet solution for AI clouds.

KEY FEATURES:

- Nearly Perfect Bandwidth at Scale
- Extremely Low Latency
- End-to-End Stack Optimisation
- Advanced RoCE Extensions for Scalable AI Communications
- Deterministic Performance and Performance Isolation
- Open Network Operating System: SONiC and Cumulus

- Improved AI Cloud performance
- Reduced TCO
- Reduced time to results
- Full-stack optimisation
- Enhanced multi-tenancy features





SOFTWARE





NVIDIA BASE COMMAND MANAGER

NVIDIA Base Command Manager is a centralised GPU resource management platform that simplifies the allocation, deployment and utilisation of GPU resources across your HPC cluster.



WAREWULF

Warewulf empowers you to create, manage and deploy system images within your HPC cluster. This flexibility enhances your cluster's efficiency by customising node configurations to meet specific workload requirements.

XCAT (EXTREME CLOUD ADMINISTRATION TOOLKIT)



xCAT is a modular and scalable HPC management solution, designed to simplify system deployment, monitoring and management across your cluster. With its bare-metal provisioning capabilities, xCAT accelerates the deployment of your HPC infrastructure.



NVIDIA BASE COMMAND MANAGER



NVIDIA Base Command Manager is the cornerstone of efficient GPU resource management for your HPC cluster. This powerful platform centralises the allocation, deployment and utilisation of GPU resources, eliminating the complexities associated with GPU management. Whether you're a seasoned data scientist or a newcomer to the world of GPU computing, NVIDIA's Base Command Manager simplifies GPU workflow, making it a valuable asset for your HPC infrastructure.

The system's user-friendly interface and robust support for multiple users provide the foundation for optimising GPU-driven workloads.

KEY FEATURES:

- NVIDIA Base Command Manager offers a unified platform to manage your GPU resources across your HPC cluster
- Facilitate collaboration with support for multiple users and seamless allocation of GPU resources
- Utilise containers to streamline Al, ML and HPC workloads, ensuring efficient resource utilisation

- Bolstered efficiency
- Easy collaboration
- Simplifies workload management



WAREWULF



Warewulf is your gateway to customised system management within your HPC cluster. This versatile tool empowers you to create, manage and deploy system images tailored to your unique requirements.

HPC clusters often handle a diverse range of workloads, each with specific node configurations and resource needs. Warewulf simplifies the process by allowing you to craft custom system images that match the exact requirements of your workloads. In the ever-shifting world of HPC, flexibility is the key to success, and Warewulf offers precisely that. It's time to streamline your cluster configuration and improve your HPC cluster's efficiency.

KEY FEATURES:

- Warewulf enables the creation and management of system images for simplified cluster configuration
- Tailor node configuration to specific workload requirements, optimising performance
- User-friendly tools simplify node provisioning and system administration

- Improved efficiency
- Customisable
- User friendly





XCAT (EXTREME CLOUD ADMINISTRATION TOOLKIT)



xCAT, or Extreme Cloud Administration Toolkit, is your comprehensive solution for the management of your HPC cluster. As HPC systems grow in both complexity and size, their administration can become increasingly challenging. xCAT, a modular and scalable system, offers a simplified way to deploy, monitor and control your overarching HPC infrastructure.

The tool's centralised management capabilities allow you to view your entire cluster through on single pane of glass and brings together system administration tasks. xCAT's bare-metal provisioning is extraordinary, allowing you to deploy systems faster and more efficiently than ever before. With xCAT, you can adapt to changes in your cluster's size and configuration, and make the most of your HPC resources.

KEY FEATURES:

- xCAT's modular architecture supports scalability and adaptability to cluster changes
- Centralise system management, making it easier to deploy, monitor and manage clusters
- xCAT simplifies the provisioning of bare-metal systems, reducing deployment times

- Scalable
- Streamlined administration
- Faster deployment





BOSTON SERVICES



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 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 LABS

Boston provides tech facilities to test emerging technologies and enterprise-grade systems from leading manufacturers. Boston Labs is customer-centric, customised solutions with off-the-shelf components to enhance business outcomes.



PROFESSIONAL SERVICES

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.



SUPPORT AND ONSITE WARRANTY

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

BOSTON TECHNOLOGY CONSULTING



Boston Technology Consulting (BTC) helps organisations to effectively create, optimise and grow their business in today's data and IT focused market. HPC consultants within BTC play a pivotal role in advising businesses on the design, implementation, and management of advanced computing systems, ensuring they are tailored to specific needs to optimize performance. Our team guide clients in selecting appropriate hardware, optimizing software, and leveraging parallel processing techniques to enhance computational efficiency. Moreover, HPC consultants stay abreast of cuttingedge technologies, emerging trends, and industry best practices, enabling organizations to stay competitive in rapidly evolving fields such as scientific research, data analytics, and AI.

46

BTC with its historical technology partnerships which stretch back over 30 years and global reach equips us with the unique skill to provide agile and scalable solutions to everyone from the smallest players to industry heavy hitters. Al is currently driving a new technology inflexion point. This is where we are finding the traditional world of Enterprise IT and HPC coming together meaning that HPC consultants are invaluable as they can leverage their knowledge to design the correct platform using the correct technologies to support Al workloads.

BOST N TRAINING ACADEMY



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 and corporate learning and consulting on Al 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.



BASTONLABS



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.

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

At Boston Labs we can give you the unique opportunity to have a handson 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.

Our line of readily available cutting-edge hardware and software 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 all 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 our findings can be read in our blog posts.



PROFESSIONAL SERVICES



49

PRESALES CONSULTATION AND SERVICES

Our Presales Consultation and Services aim to provide optimal solutions for your business needs. Our team of experts is dedicated to understanding your requirements, offering personalised consultations, and delivering our wide range of technology solutions. We go beyond product knowledge, bringing a deep understanding of industry trends and emerging technologies to help you stay ahead of the competition.

WE CAN PROVIDE:

- Consultancy and Solution Design
- Inventory and Site Survey
- Technical Account Management
- Boston Labs and POC Services
- Bespoke System Design and OEM Manufacturing
- Product and Technology updates
- Environmental and Regulatory Testing

POST SALES AND IMPLEMENTATION

Our Post Sales and Implementation team ensures a smooth and successful deployment of your technology solutions. With skilled project management expertise, thorough testing, seamless integration and comprehensive training, we prioritise your success from purchase to implementation.

WE CAN PROVIDE:

- System Integration and Burn in Testing
- Racking and Installation
- Project Management
- Global delivery
- Turnkey Solutions and Rack and Roll
- Training and Education



BOSTON'S PROFESSIONAL SERVICES CONTINUED



DATACENTRE AND HOSTING SERVICES

Your business relies on a secure, reliable and high-performance infrastructure to thrive in the digital landscape. Boston offers state-of-the-art data centre and hosting solutions designed to meet your specific needs. Our industry expertise delivers robust and scalable data centre solutions, including dedicated servers, virtual private servers (VPS), cloud hosting and colocation services. With enterprise-grade infrastructure, customised solutions, 24/7 support and scalability, we empower your digital infrastructure to reach its full potential. Choose Boston's Datacentre and Hosting Services for exceptional performance and reliability.

END OF LIFE AND EXTENDED LIFE CARE

As technology rapidly evolves, effectively managing end-of-life products and systems becomes crucial. Boston Limited's End of Life and Extended Life Care services provide comprehensive solutions to navigate this critical phase of your technology lifecycle. Our expert team understands the challenges associated with retiring outdated systems and mitigating risks related to unsupported or vulnerable hardware. With our services, we ensure a seamless transition to newer technology while maximising the value of your existing investments.

WE CAN PROVIDE:

- Decommissioning, Recycling and Disposal
- Out of Warranty Repair and Refurbishment

Please contact us if you are interested in discussing any of our additional service offerings.



WE CAN PROVIDE:

- Hosting
- Leasing
- Managed Services and Cloud
- Datacentre and Rack Design
- Rack and Roll

SUPPORT AND ONSITE WARRANTY



At Boston, exceptional customer support is the foundation of our successful partnerships. Our Support and Onsite Warranty services offer peace of mind, knowing that our dedicated team of experts is ready to assist you whenever you need help.

With prompt and reliable assistance available 24/7, our experienced professionals with deep expertise across various technologies are just a phone call away. We provide onsite support to minimise downtime and ensure smooth operations, coupled with warranty coverage for eligible hardware, ensuring swift replacement or repair.

ONSITE 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, WEEKENDS & UK BANK HOLIDAYS	NO	NO	NO	NO	NO	YES





sales@boston.co.uk | www.boston.co.uk | (+44) 1727 876 100 | Boston Limited, Unit 5 Curo Park, Frogmore, St Albans, AL2 2DD, UK Austria | Australia | France | Germany | India | Taiwan | South Africa | Switzerland | UK | USA