BOSTON
Servers | Storage | Solutions

intel select solution

Boston, Ltd. Supermicro®-Based Systems – Intel® Select Solution for Simulation and Modeling

Build your supercomputing infrastructure on Boston's extensive industry and design expertise for High Performance Computing applications

When it comes to compute-intensive research areas, like high-energy particle physics, or industrial design projects that require computer-aided engineering (CAE) and other commercial applications, simulation, and modeling workloads need to run on high-performance computing (HPC) clusters. To support those applications at scale, modern HPC systems require multi-core processors, high-bandwidth fabrics, and fast storage and other broad input/output (I/O) capabilities.

Because of the complexity and variety of technologies available on the market, designing and assembling an HPC system for specific workloads can be time-consuming, requiring specific expertise. Integration and configuration of selected components can impact the performance of the solution with the reality of properly configuring the combined solution intimidating technologists.

Boston, Ltd. (**www.boston.co.uk**) not only designs, builds, and delivers large computing clusters, it offers extensive consultancy services and expertise with early access to systems to help ensure workloads will run optimally on their systems.

Non-Volatile Memory Express (NVMe) is specifically designed to deliver high bandwidth, low latency access to storage devices over a PCI-Express interface.

Boston, Ltd. – Expertise, Experience, and Service Worldwide

Boston, Ltd. is a worldwide, leading integrator of HPC clusters used for simulation and modeling. With over 25 years of experience and vast expertise in various scientific and commercial domains, Boston services customers with design, build, integration, and installation of supercomputers, plus benchmarking and application testing, tuning, and optimization, setting them apart from other integrators.

Based on Supermicro® servers and Intel® technologies, Boston builds custom HPC solutions that appear in universities and businesses in Europe, the United States, Asia, India, Russia, and Australia. Their engineers work closely with leading organizations within Formula 1, manufacturing, engineering, and genomics.

Boston offers both potential and existing customers pre- and post-sales access to their systems and labs to help ensure that their workloads run optimally on their systems.

High-Performance Enabling Ingredients

COMPUTE

Boston offers high-performance dual-processor Supermicro servers based on Intel® Xeon® Scalable processors. Intel Xeon Scalable processors, such as the Intel® Xeon® Gold 6148 processor, feature significant enhancements that benefit HPC applications, including improvements in I/O, memory, fabric integration, and Intel® Advanced Vector Instructions 512 (Intel® AVX-512). Optionally, Intel® Xeon® Platinum processors—with up to 28 cores—can be used to meet the most challenging compute demands,[1] depending on the requirements of the workload and other customer needs.

FABRIC

Intel® Omni-Path Architecture (Intel® OPA) provides 100 gigabits per second (Gbps) bandwidth and a low-latency, next-generation fabric for HPC clusters. Its 48-port switch chip delivers a 33 percent increase in density over the traditional 36-port switch ASIC historically used for InfiniBand* networking. The denser radix chip reduces the number of required switches, cutting costs and enhancing reliability. Intel OPA also reduces cabling-related costs, power consumption, space requirements, and ongoing system maintenance requirements. These advancements can lower fabric costs by up to 61 percent.[2]

INTEL® SOLID STATE DRIVE DATA CENTER FAMILY

Intel® Solid State Drive Data Center Family (Intel® SSD DC) and Intel® Optane™ SSD DC products offer responsive performance and high capacities for local scratch storage, burst buffers, and augmented parallel files systems. Intel Optane SSD DC storage is the world's most responsive data center SSD with Intel® Optane™ technology.[3] Intel SSD DC storage includes PCIe* interface with NVMe protocol options.

SELECT SOLUTION CONFIGURATION

Table 1 lists Boston's Simulation and Modeling Solution typical configuration. Hardware and software specifications are determined by the overall workload requirements and operating system.

INGREDIENT	INTEL® SELECT SOLUTION FOR SIMULATION AND MODELING CLUSTER CONFIGURATION DETAILS
PLATFORM	Supermicro® dual-socket server-based cluster
PROCESSOR	2 x Intel® Xeon® Gold 6248 processors at 2.4 to 2.5GHz, 20 cores, 150W, 40 threads (or higher number)
MEMORY	384 GB DDR4
LOCAL STORAGE	Intel® SSD DC Family for local scratch storage, preferably PCIe* for NVMe* protocol for performance Intel SSD DC Family storage to augment parallel file system storage
MESSAGING FABRIC	Intel® Omni-Path single-port PCIe* 3.0 x16 adapter with 100 Gbps
MANAGEMENT NETWORK	Integrated 1 gigabit Ethernet (GbE)
SOFTWARE	Linux* operating system Inte® Parallel Studio XE 2018 Cluster Edition Intel® Cluster Checker 2018

Table 1: Boston, Ltd. Simulation and Modeling Solution Typical Configuration

TECHNOLOGY SELECTIONS

In addition to the Intel® Xeon® processor-based hardware foundation and Intel OPA, other technologies from Intel and Boston provide further performance gains:

Intel® AVX-512: Boosts performance for the most demanding computational workloads, with up to double the number of floating point operations per second (FLOPS) per clock cycle, compared to previous-generation Intel® processors.

Intel® Cluster Checker: Inspects more than 100 characteristics related to cluster health. Intel Cluster Checker examines the system at both the node and cluster level, making sure all components work together to deliver optimal performance. It assesses firmware, kernel, storage, and network settings. It also conducts high-level tests of node and network performance using the Intel® MPI Library benchmarks, STREAM,* the High-Performance LINPACK* (HPL*) benchmark, the High Performance Conjugate Gradients* (HPCG*) benchmark, and other benchmarks. Intel Cluster Checker can be extended with custom tests, and its functionality can be embedded into other software.

Intel® Parallel Studio: This comprehensive suite of development tools makes it simpler to build and modernize code with the latest techniques in vectorization, multithreading, multinode parallelization, and memory optimization. It enables C, C++, Fortran, and Python* software developers to:

- Increase performance: Boost application performance that scales on current and future Intel® platforms.
- Build code faster: Simplify the process of creating parallel code.
- Get Priority Support: Connect directly to Intel engineers for confidential answers to technical questions.

Intel® Cluster Runtimes: Supplies key software runtime elements that are required on each cluster to ensure optimal performance paths for applications. Intel runtime performance libraries, including Intel® Math Kernel Library (Intel® MKL) and Intel MPI Library, deliver excellent performance optimized for clusters based on Intel® architecture.

Supermicro servers feature up to four DP nodes, delivering 224 cores per 2U height with Intel Xeon Platinum Processor Family.

Supermicro storage servers offer SSD with NVMe performance and optional Intel Optane SSDs.

Boston's supercomputing solutions offer high-performance systems with worldwide professional services for demanding MPI-based simulation and modeling applications. For more information, visit Boston at **https://www.boston.co.uk**.

For more information on Intel Select Solutions, visit **intel.com/selectsolutions**.

BOSTON
Servers | Storage | Solutions

intel select solution

---

[1] Intel. "Performance Benchmarks and Configuration Details for Intel® Xeon® Scalable Processors." **intel.com/content/www/us/en/benchmarks/xeon-scalable-benchmark.html**.

[2] Configuration assumes a 750-node cluster, and the number of switch chips required is based on a full bisection bandwidth (FBB) fat-tree configuration. Intel® Omni-Path Architecture uses one fully populated 768-port director switch; the Mellanox EDR* solution uses a combination of 648-port director switches and 36-port edge switches. Intel and Mellanox component pricing is from kernelsoftware.com, with prices as of May 2016. Compute node pricing is based on the Dell PowerEdge R730* server from dell.com, with prices as of November 2015. Intel Omni-Path Architec¬ture pricing is based on estimated reseller pricing, which is based on projected Intel manufacturer's suggested retail price (MSRP) pricing at time of launch.

[3] Responsiveness defined as average read latency measured at queue depth 1 during 4k random write workload. Measured using FIO 2.15. Common configuration – Intel® 2U Server System, OS CentOS* 7.2, kernel 3.10.0-327.el7.x86_64, CPU 2 x Intel® Xeon® E5-2699 v4 @ 2.20GHz (22 cores), RAM 396 GB DDR @ 2133MHz. Intel drives evaluated – Intel® Optane™ SSD DC P4800X 375GB and Intel® SSD DC P3700 1600 GB. Samsung drives evaluated – Samsung SSD* PM1725a, Samsung SSD* PM1725, Samsung* PM963, Samsung* PM953. Micron drive evaluated – Micron* 9100 PCIe* NVMe* SSD. Toshiba drives evaluated – Toshiba* ZD6300. Test – QD1 Random Read 4K latency, QD1 Random RW 4K 70% Read latency, QD1 Random Write 4K latency using FIO 2.15.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit intel.com/benchmarks.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

For more complete information about performance and benchmark results, visit intel.com/benchmarks.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

0618/KM/HB/PDF Please Recycle 337664-001-01US