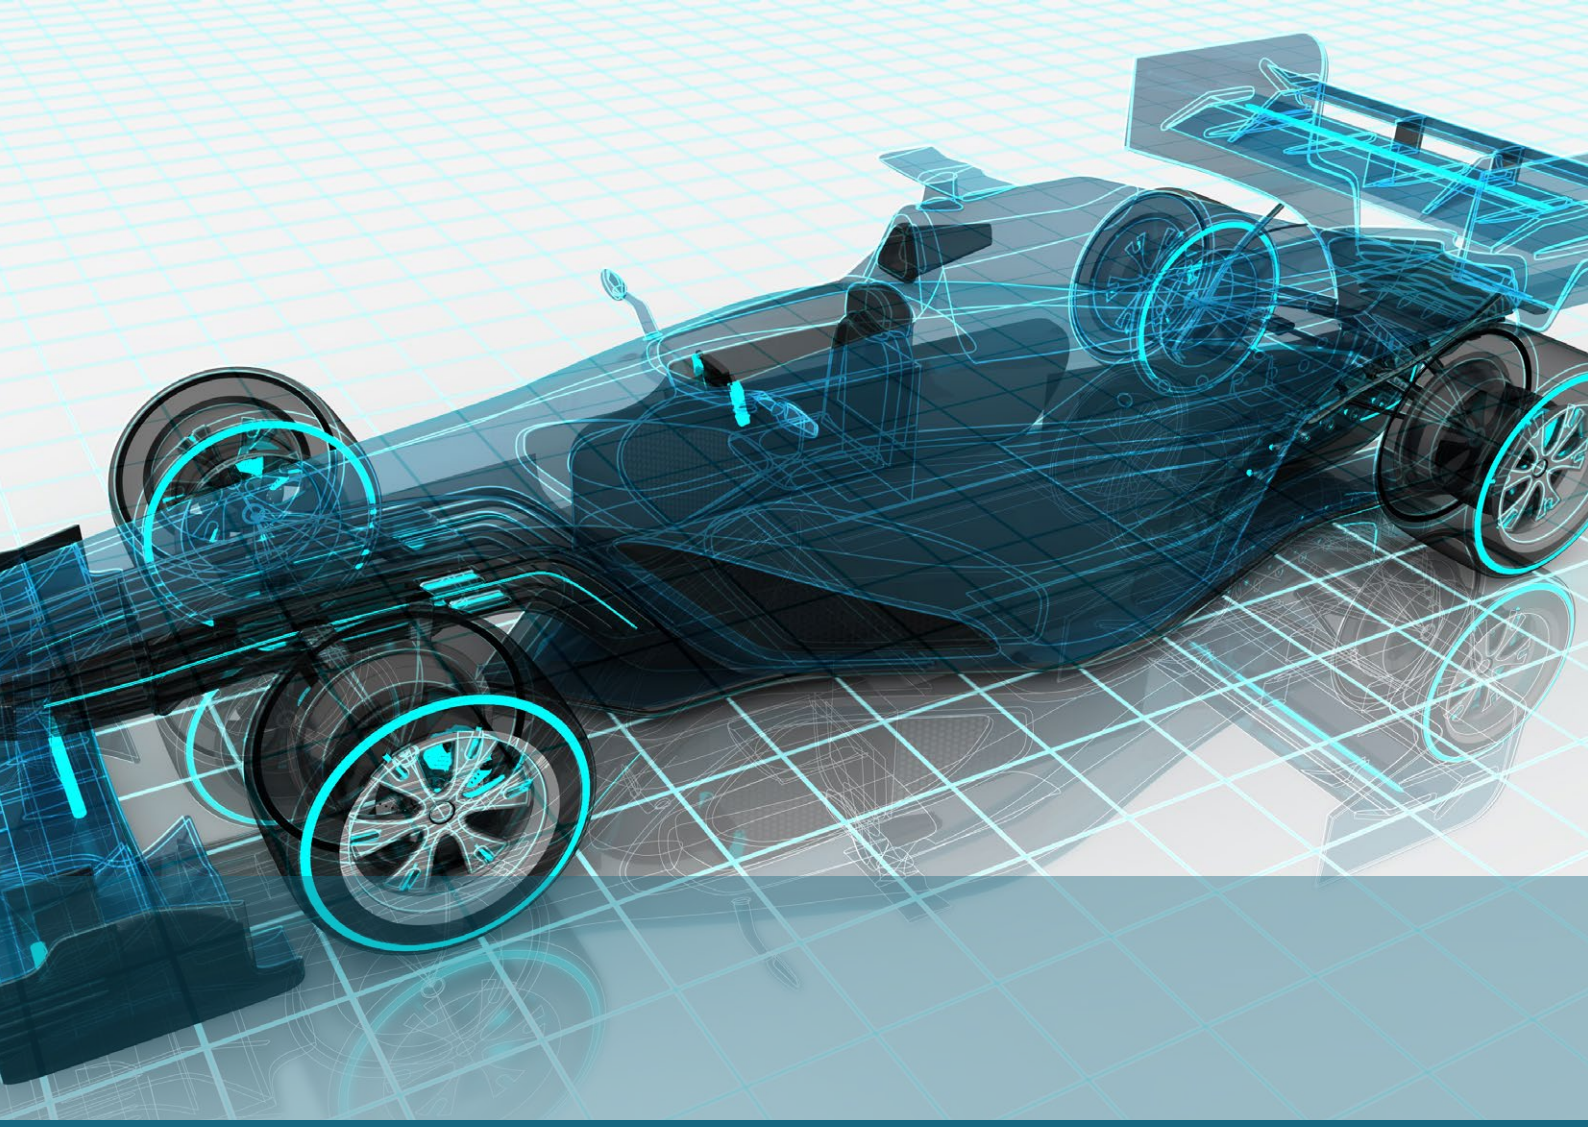

Overcoming challenges of HPC Clusters for Simulation & Modeling

Produced by Europa Science in conjunction with:

SCIENTIFIC COMPUTING WORLD

BOSTON
Servers | Storage | Solutions

BOSTON
Servers I Storage I Solutions

High-performance computing (HPC) based simulation and modeling has become a fundamental part of engineering design and research. Whether organizations are facing increasing competition or want to accelerate time to market for new products, HPC provides a platform to accelerate simulation workloads. Engineers working in an increasingly broad range of industries from traditionally computational demanding disciplines such as aerospace and automotive, to more recent adopters including construction and architecture, robotics, energy production, and healthcare are using HPC to execute simulation workloads.

These companies are seeking to employ HPC to solve engineering challenges in Computational Fluid Dynamics (CFD), to enhance multiphysics simulations or improve structural optimization of components. They want to design, analyze and certify components and systems using simulation. This enables them to reduce physical testing and verification, increase the number of design iterations, or develop more complex products in shorter development cycles. To solve these challenges they need the power of supercomputing.

HPC provides a platform to accelerate simulation workloads

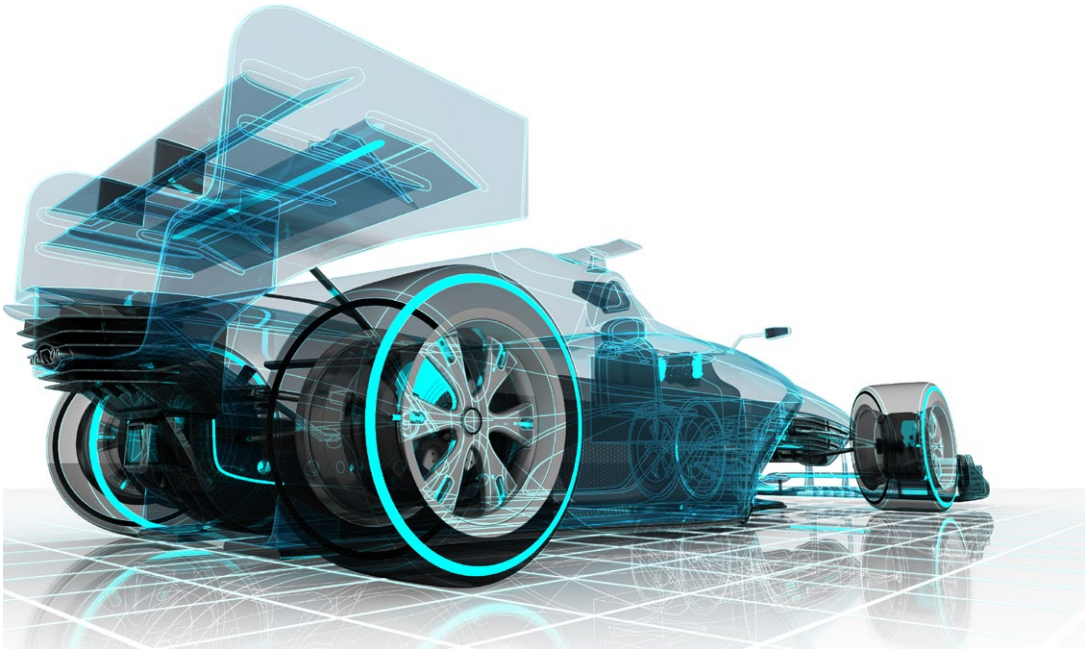

BOSTON
Servers | Storage | Solutions

Optimised for modeling and simulation

Simulation industry verticals such as automotive and aerospace represent some of the early adopters of HPC technology. With complex and large scale product development programs, organizations in these industries have found that HPC can reduce the amount of physical testing required to validate new designs - drastically reducing project timelines. If an aerodynamic simulation of a new vehicle can be conducted in just a few hours, instead of days or weeks, engineers can make agile decisions and drive performance improvements. HPC can scale simulation capabilities to meet new design challenges and increase industry competitiveness.

HPC can scale simulation capabilities to meet new design challenges and increase industry competitiveness

Whether designing a new advanced driver-assistance systems (ADAS) system or the latest F1 engine, HPC enables automotive engineers to continuously drive product innovation.

As the use of simulation has grown, software packages, including computer-aided engineering (CAE) and CFD applications have grown and been adapted to HPC clusters. But, to run these applications at scale, HPC systems require multi-core processors, a high-bandwidth fabric, and sufficient storage and I/O performance to provide a balanced compute infrastructure that is not bottlenecked by hardware.

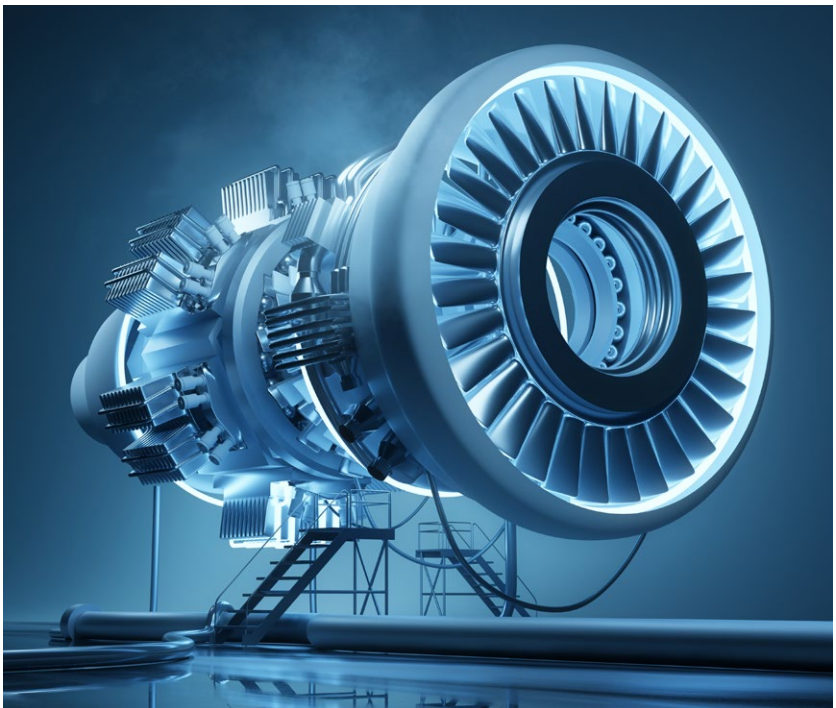



BOSTON
Servers I Storage I Solutions

The Boston Intel® Select Solution provides an easily accessible path to a proven HPC architecture for CFD and other simulation workloads. This technology has been specifically designed to overcome challenges with hardware selection and software performance to ensure simulations run optimally.
Build your supercomputing infrastructure with the Boston Intel Select Solution for Simulation and Modeling - a high-performance solution optimized for demanding simulation workloads. Boston is now offering pre-validated and tested solutions that combine 1st and 2nd Generation Intel Xeon Scalable processors in a proven architecture that delivers guaranteed performance for modeling and simulation workloads.

The Boston Intel Select Solution enables organizations to scale quickly with a workload-optimized solution that contains the latest HPC hardware including multicore processors, a high bandwidth fabric and fast storage. Boston provides end-to-end support from rack-scale design to networking, cabling and installation, removing the challenges traditionally associated with procuring and managing the setup of a new HPC system.

Users can also take the opportunity to test their own workloads on the Intel Select Solution at Boston's research and development facilities. Engineers can test their real-world workloads to determine the performance they can achieve with their own real-world CFD or simulation codes.

Boston provides end-to-end support from rack-scale design to networking, cabling and installation

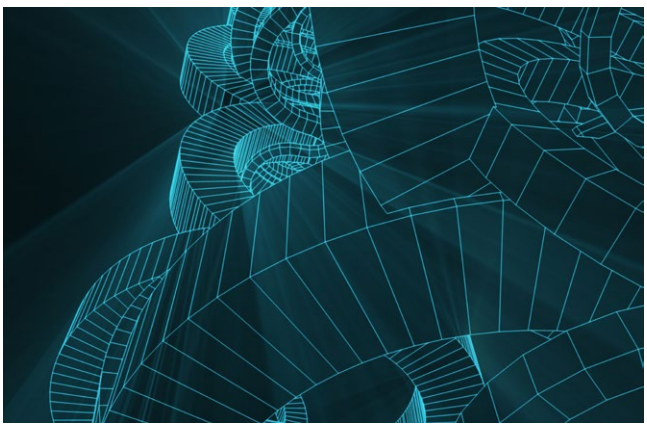



BOSTON
Servers I Storage I Solutions

Pre-configured, proven performance

The Boston Intel Select Solution comes in both Base and Plus configurations which have been designed specifically to offer a simple path and quick-to-deploy infrastructure. The solutions and hardware choices are selected using a standards-based approach defined in the Intel HPC Platform Specification. This provides peace of mind to engineers and researchers by ensuring interoperability with common applications used in modeling and simulation today.

All Intel Select Solution products pre-validated and tested to ensure that key components are used which are compliant with industry standards and best practices for Intel-based clusters. These solutions also must meet or exceed defined performance levels in targeted characteristics important to HPC applications.

Base and Plus configurations of Intel Select Solutions for Simulation & Modeling[1]

Ingredient	Intel Select Solutions for Simulation & Modeling Base Configuration	Intel Select Solutions for Simulation & Modeling Plus Configuration
Workload Domain (Minimum 4-Node Configuration)		
Platform	Dual-socket server platform	Dual-socket server platform
Processor	2 × Intel® Xeon® Gold 6126 processor (2.60 GHz, 12 cores/24 threads), Intel® Xeon® Gold 6226 processor (2.7 GHz, 12 cores/24 threads), Intel® Xeon® Gold 6226R processor (2.90 GHz/16 cores/32 threads), or a higher model number Intel® Xeon® Scalable processor	2 × Intel® Xeon® Gold 6148 processor (2.40 GHz, 20 cores/40 threads), Intel® Xeon® Gold 6252 processor (2.5 GHz, 12 cores/24 threads), Intel® Xeon® Gold 6248R processor (3.00 GHz, 24 cores, 48 threads), or a higher model number Intel Xeon Scalable processor
Memory	96 GB (12 × 8 GB 2,666-MHz 288-pin DDR4 RDIMM) 2 GB memory per processor core and all memory channels populated	96 GB (12 × 8 GB 2,666-MHz 288-pin DDR4 RDIMM) 2 GB memory per processor core and all memory channels populated
Local Storage	1 × Intel® SSD DC S3520 Series or better, or Intel SSD DC P3520 Series or better**	1 × Intel SSD DC S3520 Series or better, or Intel SSD DC P3520 Series or better**
Messaging Fabric	1 × Intel® Omni-Path Architecture (Intel® OPA), single-port Peripheral Component Interconnect Express (PCIe) 3.0 x16 adapter, 100 gigabits per second (Gbps)	1 × Intel® OPA, single-port PCIe 3.0 x16 adapter, 100 Gbps

Intel



BOSTON
Servers | Storage | Solutions

Compute:
Processor performance is delivered in Supermicro servers based on Intel Xeon Scalable processors including the Gold 6148 which offers 20 cores to deliver exceptional performance for compute and data-intensive workloads. Users can also opt for Intel Xeon Platinum processors—with up to 28 Cores.

2nd Generation Intel Xeon Scalable processors feature significant enhancements that can benefit HPC applications, including improvements in input/output (I/O), memory, and Intel Advanced Vector Extensions 512 (Intel AVX-512).

For HPC users looking to adopt AI, Intel Deep Learning Boost (Intel DL Boost) uses the Vector Neural Network Instructions (VNNI) set to accelerate the performance of AI deep learning (inference) workloads, such as speech recognition, image recognition, object classification, and machine translation.

2nd Generation Intel Xeon Scalable processors feature significant enhancements that can benefit HPC applications

Fabric:
The Intel Omni-Path Architecture (Intel OPA) provides 100 gigabits per second (Gbps) bandwidth and a low-latency, next-generation fabric for HPC clusters. Intel OPA can also reduce cabling-related costs, power consumption, space requirements, and ongoing system-maintenance requirements.

Storage:
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 Select Solution.

Configuration:
Hardware and software specifications are determined by the overall workload requirements and operating system.



The Intel Select Solution is based on several underlying technologies which further support hardware and software performance:

Intel® AVX-512:
Boosts performance for the most demanding workloads with up to 2x FLOPS per clock cycle vs. predecessor

Intel® Cluster Checker:
Examines the system at both the node and cluster level, ensuring components work together for optimal performance

Intel® Parallel Studio:
This suite of development tools makes it simple to build and modernize codes, enabling C, C++, and Python developers to increase performance, build code faster, and get priority support

Intel® Cluster Runtimes:
Supplies key software runtime elements that are required on each cluster to ensure optimal performance paths for applications

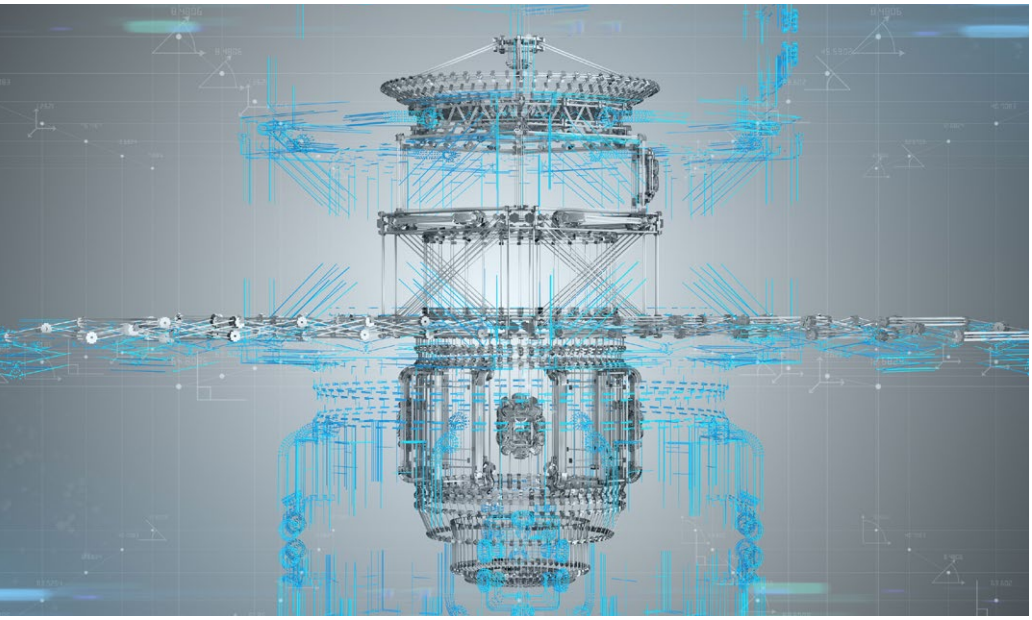



BOSTON
Servers I Storage I Solutions

Benchmark and user testing

All Intel Select Solutions are verified to meet a specified minimum level of workload-optimized performance capabilities. The solutions must meet or exceed design and testing standards across five well-known industry benchmarks that cover important system aspects and indicate potential scale-up and scale-out performance for simulation and modeling application workloads.

Intel currently benchmarks performance using DGEMM, STREAM and IMB PingPong to measure the performance of key characteristics of the system: compute power, memory bandwidth, and interconnect fabric performance. Additionally, two popular benchmarks are used to represent applications, The High Performance LINPACK (HPL) and the High Performance Conjugate Gradient (HPCG) benchmark.

One of the major benefits of selecting the Boston Intel Select Solution is that these systems provide a guaranteed level of performance to users

One of the major benefits of selecting the Boston Intel Select Solution is that these systems provide a guaranteed level of performance to users. However, Boston offers an additional level of service by enabling potential users to test their own application codes on the Intel Select Solution at Boston's research and development facilities before purchase. This gives organizations access to test the latest Intel technologies, in some cases, before they have been officially announced. Users can test their applications in a real-world setting using the same hardware they will receive to demonstrate real-world performance.

Boston Limited has been providing cutting-edge technology since 1992 using Supermicro building blocks—their high performance, mission-critical server and storage solutions can be tailored for each client, helping customers to create the ideal solution for their modeling and simulation requirements.

For more information about Boston or the Boston Intel Select Solution please contact us with your questions.

Reference:

[1] www.intel.co.uk/content/www/uk/en/products/solutions/select-solutions/hpc/simulation-modeling.html