WHITE PAPER

Manufacturing
Digital Design and Engineering

Higher Cluster Performance for ANSYS® Fluent® with Intel® Omni-Path Architecture

Intel® Omni-Path Architecture enabled from 25 to 47 percent higher performance versus EDR InfiniBand* across a range of ANSYS® Fluent® simulations run on large cluster configurations.[1]

"Thanks to Intel® Omni-Path Architecture and the latest Intel® Xeon® processor E5-2600 v4 product family, ANSYS® Fluent® is able to achieve performance levels beyond our expectations. Its unrivaled performance enables our customers to simulate higher-fidelity models without having to expand their cluster nodes."

Dr. Wim Slagter, Director of HPC and cloud marketing, ANSYS

Many engineering design teams are forced to make trade-offs in their digital simulations, sacrificing some amount of detail and accuracy in exchange for fast runtimes that fit into tight design schedules. Such trade-offs become even more problematic for large simulations that involve not only moving parts, but also other complex phenomena, such as turbulence, chemical reactions, radiation, or heat transfer.

ANSYS and Intel help organizations address this challenge through optimized solutions that can scale efficiently across large numbers of servers to deliver desired performance levels. One example is ANSYS® Fluent® running on the Intel® Xeon® processor E5 v4 family. This hardware and software combination scales up to thousands of Intel® Xeon® processor cores to provide fast turnarounds for large, complex simulations.

Intel® Omni-Path Architecture (Intel® OPA) adds to these advantages, providing a high-speed cluster fabric that helps to unleash the full power of the processors in clustered configurations. This high-performance, low latency fabric is designed to resolve the performance, scalability, and cost challenges of traditional InfiniBand* solutions so organizations can get higher value from their high performance computing (HPC) investments.

Superior Performance Scaling for ANSYS® Fluent® with Intel® Omni-Path Architecture

Combustor_12m, ANSYS Fluent (Intel® Xeon® processors E5-2697 v4)

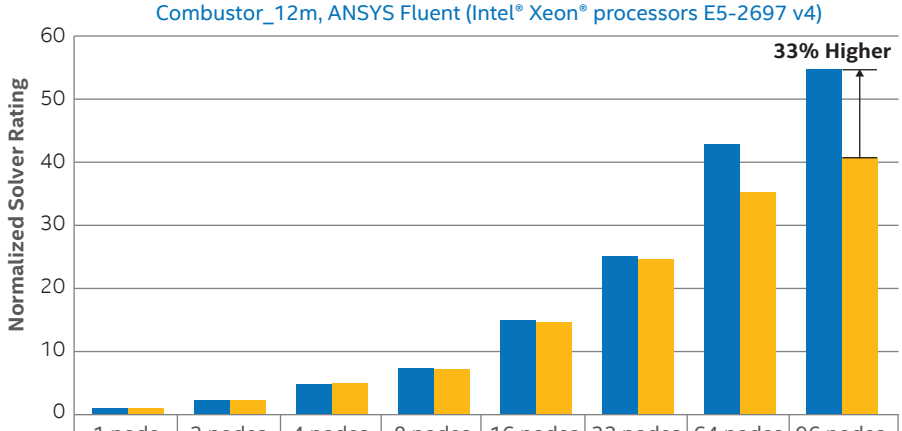

	1 node	2 nodes	4 nodes	8 nodes	16 nodes	32 nodes	64 nodes	96 nodes
Intel® Omni-Path Architecture	1	2.00	3.95	7.63	14.45	24.51	42.37	54.57
EDR InfiniBand*	1	2.00	3.99	7.60	14.40	24.00	35.50	41.00

Figure 1. Intel® Omni-Path Architecture helps to improve performance scaling for ANSYS® Fluent® compared with EDR InfiniBand*, enabling up to 33 percent higher performance for this 12-million cell model running on a 96-node cluster, with comparable gains across a range of models.[1]

Faster Time to Results on Medium to Large Clusters

ANSYS and Intel measured performance for ANSYS Fluent on Intel Xeon processors for a diverse set of models that ranged in size from 2 million to 14 million cells. Simulations were run on single servers and on progressively larger clusters up to a maximum of 96 nodes (3,456 cores). Two different cluster interconnect fabrics were used: Intel® OPA and EDR InfiniBand*.

For every model, the Intel® OPA cluster enabled performance that was either closely comparable or superior to EDR InfiniBand*. In all test scenarios, the scalability advantages provided by Intel® OPA were most pronounced at large cluster sizes, with gains ranging from 25 to 47 percent on the largest tested cluster for each model (Figure 2).

Digital Model	Relative Performance (Intel® Omni-Path Architecture/EDR InfiniBand*) Number of Clustered Servers (cores)								Performance Gain (using Intel® OPA on the largest tested cluster size)
	1 (36)	2 (72)	4 (144)	8 (288)	16 (576)	32 (1,152)	64 (2,304)	96 (3,456)	
Fluidized_bed_2m	1.00	1.00	0.98	1.04	1.06	1.25	—	—	25%
Pump_2m	1.00	1.0	0.98	1.02	1.04	1.16	1.30	—	30%
Rotor_3m	1.00	0.99	1.00	1.03	1.05	1.17	1.47	—	47%
Sedan_4m	1.00	1.00	0.99	1.00	0.98	1.14	1.39	—	39%
Combustor_12m	1.00	1.00	0.98	1.00	1.00	1.02	1.19	1.33	33%
Aircraft_wing_14m	1.00	0.99	0.99	0.99	1.00	0.94	1.11	1.25	25%

Figure 2. Effect of Cluster Fabric on ANSYS® Fluent® Performance[1]

A Fabric for the Future of HPC

Superior scalability is just one of the benefits provided by Intel® OPA. This interconnect fabric was designed specifically to address the performance and cost challenges that organizations face when deploying traditional fabric solutions. Intel® OPA transmits data with the same 100 Gbps line speed provided by EDR InfiniBand*, but this tells only part of the story. Innovative new features help to improve performance, efficiency, resilience, and cost models across a wide range of HPC workloads.

- **Low-latency even at extreme scale**. Intel® OPA uses streamlined, lightweight PSM2 libraries that help to reduce latency versus the VERBS-based software used in InfiniBand* (VERBS was originally written for storage connectivity and includes millions of lines of code). The streamlined software helps to reduce communications overhead to enable low latency and high message rates at scale. It is particularly beneficial for applications, such as ANSYS Fluent, that generate high volumes of small MPI messages.
- **Improved fabric efficiency**. Intel® OPA provides advanced traffic shaping to dynamically optimize data flow and fabric utilization. It also includes quality-of-service features that help to provide predictable low latency for MPI messages in mixed traffic environments. This allows storage and MPI traffic to share the same Intel® OPA fabric, without slowing performance for critical applications.
- **Improved resilience**. Unlike InfiniBand*, Intel® OPA implements no-latency error checking on every link to improve data accuracy without slowing performance. When errors are identified, retries can be sent immediately across the individual link, without the long delays generated by end-to-end retries. Intel® OPA also stays up and running in the event of a physical link failure, so time-critical simulations can run to completion.
- **Better economics**. Intel® OPA switches are based on a 48-port chip architecture versus 36-port for InfiniBand*. The higher port density provides both performance and cost advantages in medium to large clusters, by reducing the number of switches, cables, and switch hops. The reduction in fabric costs can often enable organizations to purchase and connect more servers within the same budget to extend the value of their HPC investments.[2]

Ongoing Advances through Collaborative Innovation

Intel Xeon processors and Intel® OPA are both components of the Intel® Scalable System Framework (Intel® SSF). That means they have been developed and tested together—and with other components of the HPC solution stack—to deliver optimized performance and higher value in HPC environments (see the sidebar, Transforming the Future of HPC).

ANSYS and Intel are working together to ensure that ANSYS software is optimized to take full advantage of the Intel SSF solution stack, including compute, memory, fabric, storage, and software solutions. ANSYS Fluent 18.0, for example, includes the Intel® MPI Library for highly optimized message passing. It also includes full support for Intel® OPA to simplify the creation of scalable, high-performing clusters.

Next-generation ANSYS Fluent 18.1 will drive performance and efficiency even higher by adding optimized support for the most recent generation of Intel® Xeon Phi™ processors. The advanced parallelism of these processors (up to 72 cores and 288 threads per socket) will help to boost per-node performance for the largest and most complex models. These processors are also available with integrated high-speed memory and integrated Intel® OPA controllers to help improve scaling and reduce total costs in clustered environments. With ANSYS Fluent 18.1, Intel Xeon Phi processors, and Intel® OPA, many organizations will be able to achieve powerful performance benefits out of the box for some of their most demanding simulations

Take the Next Step

ANSYS and Intel are driving innovations that can help engineering teams extract deeper and more accurate insights from their digital simulations, while reducing runtimes and improving HPC cost models. Visit the following websites for more information.

- **ANSYS Fluent**. http://ansys.com/Products/Fluids/ANSYS-Fluent
- **Intel Scalable System Framework**. http://www.intel.com/content/www/us/en/high-performance-computing/product-solutions.html
- **Intel® Omni-Path Architecture.** http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/overview.html

Transforming the Future of HPC

Traditional HPC is hitting a performance wall as models become more complex and data volumes grow at exponential rates. Intel is delivering innovation and tighter integration across the complete HPC solution stack to break down those walls.

Intel innovations in memory, fabric, storage, and software help to ensure that system-level performance keeps pace with ongoing gains in processing power. A key strategy is to move everything closer to the processors to improve bandwidth, reduce latency, and drive down costs. The goal is to provide a far more scalable foundation for growing HPC requirements, while dramatically raising the bar for performance and value at every scale.

Learn more from:

- ANSYS: http://www.ansys.com/Support/resource-library/video/hpc-simplified
- Intel: http://www.intel.com/content/www/us/en/high-performance-computing/product-solutions.html

[1] T Intel performance tests conducted October 2016 using ANSYS® Fluent® 17.2 (ANSYS 18.0 is based on the same core software code and can be expected to deliver similar or superior performance in most scenarios). Tests were performed on 1 to 96 server nodes using Intel® Omni-Path Architecture (Intel Corporation Device 24f0 — Series 100) versus EDR InfiniBand* (Mellanox EDR MT4115 ConnectX-4, F/W version 12.16.1020). Server node configuration: 2 X Intel® Xeon® processor E5-2697 v4 (18 core, 2.3 GHz; 36 total cores per server), 128 GB DDR4 memory @ 2400 MHz (8 x 16 GB DIMMS), Red Hat Enterprise Linux* 7.2.

[2] Based an internal Intel analysis comparing 750-node fabrics based on Intel Omni-Path Architecture (Intel® OPA) and EDR InfiniBand* and a 948-node fabric based on Intel® OPA. Results showed that Intel® OPA could enable 26 percent more servers to be purchased and connected within the same cluster budget. EDR InfiniBand* costs and capabilities were based on Mellanox product specifications and costs available online as of September 24, 2015. **750-node EDR InfiniBand* fabric configuration**: 2 x MCS7500 plus 648-EDR chassis switch; 22 x MCS7510-E 36-port EDR 100 Gb/s InfiniBand* leaf blades; 42 x MSB7790-ES2F 36-port switches; 750 MCX455A-ECAT ConectX-4 VPI adapter cards; 750 passive copper cables; 378 active fiber cables; total estimated fabric cost of USD 3,092,459. **750-node Intel® OPA fabric configuration**: 1 x Intel® OPA 100 Series DCS (24-slot); 24 x Intel® OPA 100 Series 32-port Leaf; 750 x Intel® OPA 100 Series Host Fabric Interface; 188 passive copper cables, 562 active fiber cables; total estimated fabric cost of USD 1,217,282. 948-node Intel® OPA fabric configuration: 1 x Intel® OPA 100 Series Edge Switch (managed); 63 x Intel® OPA 100 Series Edge Switch (unmanaged); 948 x Intel® OPA 100 Series Host Fabric Interface; 948 passive copper cables, 960 active fiber cables; total estimated fabric cost of USD 1,718,655. Compute nodes in all fabric scenarios are 2-socket Intel® Xeon® processor-based servers at an estimated USD 9,000 per server. For more information, read the Intel solution brief: http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/high-performance-low-cost-brief.html#

 0814/JRR/MESH/PDF 328678-005US