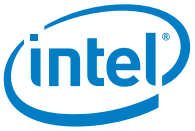

White Paper
Clifford Oberholtzer
Product Marketing – Fabric Group
Intel Corporation

ANSYS® Fluent® Brings CFD Performance with Intel® Processors and Fabrics

A Generational Performance Study

<April 9, 2015>

Executive Summary

It takes a convergence of numerous factors in a High Performance Computing (HPC) cluster to drive application performance as both computer power and software efficiency scale. The main goal is for the CPU/memory complex along with the HPC Fabric interconnect to provide scalable computing power to the application. Intel® Xeon® processors and the Intel® True Scale Fabric provides the scalable power used by ANSYS® Fluent® software and many other HPC applications. Performance results demonstrate how effectively performance can scale when a purpose designed HPC cluster network such as the Intel® True Scale Fabric, supplies the cluster scaling power. All while each generation of application software becomes more efficient and each CPU generation gains in processing power.

HPC applications need to scale in every aspect of clusters to reduce the wall-clock time needed get to the required answers. ANSYS Fluent is a state-of-the-art application that falls into the Computational Fluid Dynamics (CFD) classification. It contains the wide-ranging physical modeling capabilities needed to model flow, turbulence, heat transfer etc. An example of this is the ability to model air flow over a plane wing, before it is even built. This reduces overall design costs and enables design optimization before the first prototype needs to be built.

Performance results demonstrate how effectively performance can scale when a purpose designed HPC cluster network such as the Intel® True Scale Fabric, supplies the cluster scaling power.

The generational performance in this HPC cluster is driven by 3 major components:

ANSYS Fluent

ANSYS Fluent software contains the broad physical modeling capabilities needed to model flow, turbulence, heat transfer, and reactions for industrial applications ranging from air flow over an aircraft wing to combustion in a furnace, from bubble columns to oil platforms, from blood flow to semiconductor manufacturing, and from clean room design to wastewater treatment plants. Special models that give the software the ability to model in-cylinder combustion,

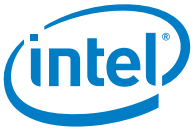

aeroacoustics, turbomachinery, and multiphase systems have served to broaden its reach.

Generational Intel® Xeon® processor comparison.

Intel® Xeon® Processor	E5-2680	E5-2697V2	E5-2697V3
Intel® Smart Cache	20 MB	30 MB	35 MB
Intel® QPI Speed	8 GT/s	8 GT/s	9.6 GT/s
# of QPI Links	2	2	2
Instruction Set	64-bit	64-bit	64-bit
Instruction Set Extensions	AVX	Intel® AVX	AVX 2.0
Lithography	32 nm	22 nm	22 nm
# of Cores	8	12	14
# of Threads	16	24	28
Processor Base Frequency	2.7 GHz	2.7 GHz	2.6 GHz
Max Turbo Frequency	3.5 GHz	3.5 GHz	3.6 GHz
Max Memory Size (dependent on memory type)	384 GB	768 GB	768 GB
Memory Types	DDR3-800/1066/1333/1600	DDR3-800/1066/1333/1600/1866	DDR4-1600/1866/2133
Max # of Memory Channels	4	4	4
Max Memory Bandwidth	51.2 GB/s	59.7 GB/s	68 GB/s

Intel® True Scale Fabric.

Intel® True Scale Fabric is an end-to-end InfiniBand implementation purpose designed from the ground up to bring high performance to MPI-based applications. The Intel® True Scale Fabric solution provides an open source host stack. The main part of this stack is Performance Scaled Messaging (PSM), a host interface is an innovation delivered as part of the Intel® solution. It utilizes connectionless traffic processing, ensuring no performance robbing cache misses and very high message rates with low end-to-end MPI latency and high effective application bandwidth. This enables MPI applications to scale from hundreds to thousands of nodes. The Intel® True Scale Fabric has sufficient performance headroom to support succeeding generations of both processor and MPI applications. The Intel® True Scale Fabric is used for all testing scenarios'.

Generational Server Performance

The Intel® True Scale Fabric is at the center of every test conducted. It brings the performance criteria required to enable applications, such a Fluent to higher and higher rating. The criteria include effective use of available bandwidth by supporting a high message rate and low end-to-end latency. While not every application takes advantage of available communications power, testing shows that in most Fluent benchmark applications the Intel® True Scale Fabric scales well with each application improvement and each Intel® Xeon® generation.

The following standard Fluent benchmark ratings were used as the performance measuring stick:

- aircraft_2m
- eddy_417k
- sedan_4m
- truck_14m
- truck_poly_14m
- truck_111m
- turbo_500k

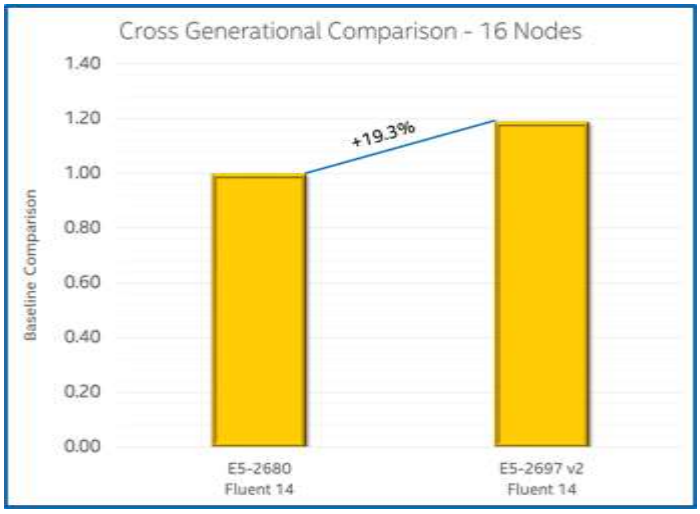

Figure 1: E5-2680 versus E5-2697 v2

Initial testing results compare processor generations using the same Fluent version 14. Dual processor servers with E5-2680[1] processors as a baseline were compared against Dual processor servers with E5-2697 v2[2] Processors.

On average the E5-2697 v2 processor-based cluster outperformed the E5-2680 processor-based cluster by over 19%.

This does not mean however, that every benchmark application received the same performance benefit from the E5-2687 v2 processor, Fluent application and HPC fabric. In fact the actual performance ratings for each benchmark varied greatly. The Solver ratings ranged from roughly equivalent to the previous processor generation to an over 46% gain.

[1] Dual 2.7Ghz E5-2680, 64GB 1333Mhz DDR3 Memory, Intel MPI, Fluent v14
[2] Dual 2.7Ghz E5-2697 v2, 64GB 1600Mhz DDR3 Memory, Intel MPI, Fluent v14

This is due to the fact that every application or benchmark has its own performance "fingerprint", its own special blend of Software, OS, processor, memory subsystem and fabric needs.

In these tests, the same high performance HPC fabric and Fluent version were used. This means that the processor with its supporting chips and memory performance and motherboard design were the major variables that changed from the previous generation. As seen in Figure 2, the Truck_Poly_14m benchmark solver rating is increased by over 40%. Being a test case utilizing 14 million cells, it fits very well to the 16 nodes used here. In addition, the segregated solver is being used, so memory bandwidth requirement is not as high as the coupled solver. That means it will take advantage of 50% more cores and the Intel® True Scale Fabric. .

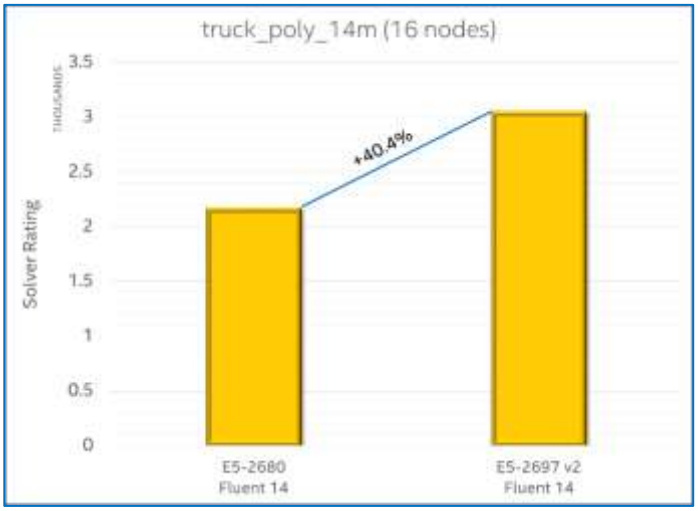

Figure 2: Fluent Truck_Poly_14m Results

On the other hand, the E5-2697 v2 server generation only managed a 4.5% gain for the Fluent Turbo_500k benchmark, as seen in Figure 3. The reason why this case is not showing much more improvement is the fact that it's a much smaller case, just half a million cells, which may be too small for 16 node runs, where it has 384 cores for E5-2697 V2 50% more than the E5-2680. Smaller problem sizes start to see diminishing returns in distributed memory parallelization sooner than larger problem sizes.

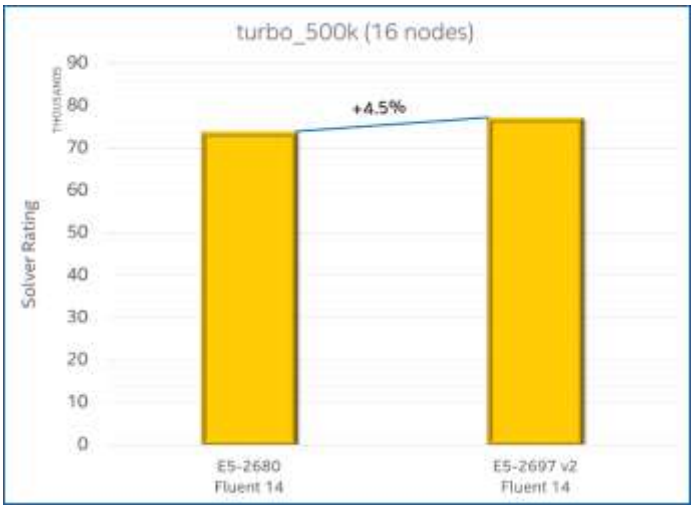

Figure 3: Fluent Turbo_500k Results

This variation in performance gains is one of the key reasons that it is recommended that customer applications be run on an actual cluster. While micro-benchmarks can give some indication of cluster performance, it does not paint the whole picture. The Intel® True Scale Fabric is available for customer testing in one of several available locations.

Generational Server and Software Performance

In these tests, not only is there a new generation of processors (E5-2697 v3)[3] with even faster memory (2133Mhz), a new generation of Fluent software (v15) was used. All results shown will demonstrate the additional solving power over the previous generation.

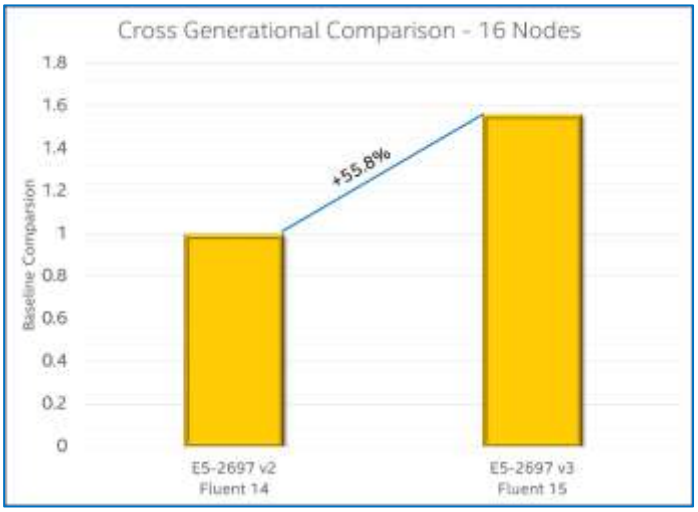

Figure 4: E5-2697 v2 versus E5-2697 v3

As seen in Figure 4 the combination of next generation processor hardware and Fluent application software coupled with the same Intel® True Scale Fabric help drive exceptional Solver Ratings. In fact overall performance surges over 55%.

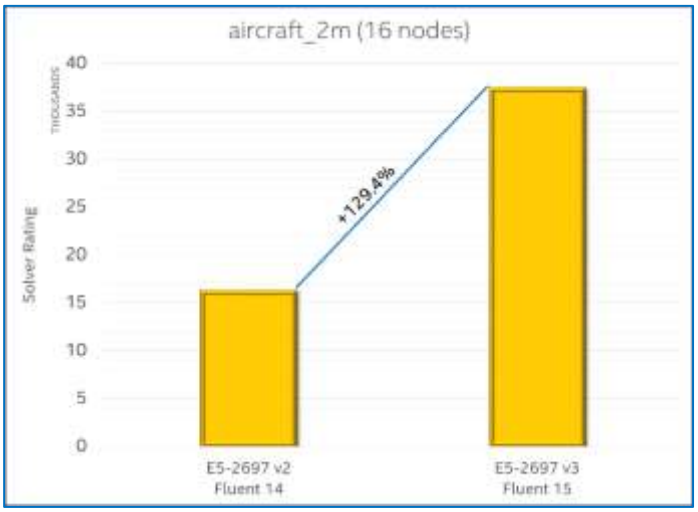

Figure 5: Fluent Aircraft_2m Results

As with the previous results each application received a different benefit from this new generation. All Solver ratings increased far more than previous results. In fact generational improvement ranged from ~12% to over 129% as shown in Figure 5.

[3] Dual 2.6Ghz E5-2697 v3, 64GB 2133Mhz DDR4 Memory, Intel MPI, Fluent v15

Conclusion

When used together the ANSYS Fluent application and Intel® Xeon® processors using the Intel® True Scale HPC Fabric enable higher generational performance. And continue to provide excellent Solver ratings at scale across the generations. When Sandy Bridge is used as a baseline, performance gains are shown through each successive processor and application generation all while using the same HPC Fabric optimized from the ground up for HPC MPI applications such as Fluent.

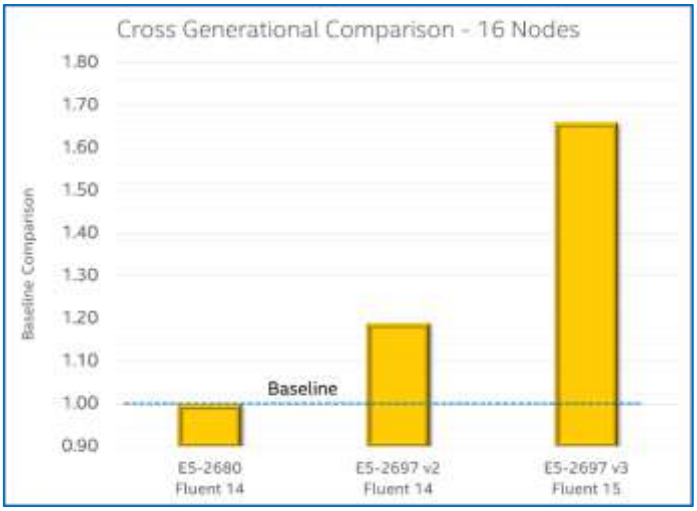

Figure 6: Generational Performance

§