

Getting the best performance out of ANSYS Fluent with Huawei high-density server for your automotive CFD simulations



It is a well known fact that high-performance computing (HPC) is vital in solving difficult computational problems across a diverse range of scientific, engineering, and business fields. HPC applications are not only being used in highly specialized industries to conduct advanced analytics and simulations, but have also been extended to the enterprise domains in various fields such as automotive, aerospace, oil and gas, financial services and the life sciences.

In the automotive industry, HPC applications are mainly used in crash simulation, computational fluid dynamics (CFD) and structural analysis. There are various parameters such as processor power, network bandwidth and latency, and memory frequency that can affect the performance of applications in HPC systems. So it is crucial to determine the exact configurations to identify the best performance at a good price.

Through this white paper, Huawei found the optimal configuration for ANSYS Fluent, a famous CFD simulation software product in various industries, by running benchmark cases on a Huawei HPC cluster with optimized system parameters.

First, we will give brief summary of ANSYS Fluent and the Huawei HPC Solution. Second, we will introduce the test environment, including benchmark cases, the Huawei X6800 server and detailed test configurations. Third, we will describe the details of test results based on solver rating and speedup in various configurations, such as single node and multi-nodes, and turbo mode and hyper-threading setting, based on network and mpi. Finally, we will give some pertinent suggestions about configuration parameters based on the conclusion of these test results.

Brief summary of ANSYS Fluent and Huawei HPC Solution

Here is a brief introduction about ANSYS Fluent and Huawei HPC Solution that were used to conduct the tests.

About ANSYS Fluent

ANSYS is the recognized leader in engineering simulation worldwide. With the most comprehensive product solution set, integrating fluids, structures, electronics, semiconductors, systems and more, ANSYS provides unparalleled multiphysics simulation.

ANSYS Fluent is the most powerful computational fluid dynamics (CFD) software available, containing the broad physical modeling capabilities needed to model flow, turbulence, heat transfer and reactions for industrial applications such as automotive manufacturing, oil platforms, semiconductor manufacturing and so on. For more details see: <http://www.ansys.com/products/fluids/ansys-fluent>

About Huawei HPC Solution

Huawei is dedicated to providing innovations on HPC infrastructures and solutions for customers to enable HPC optimization for new workload types and accelerate HPC convergence on the cloud & big data.

Huawei HPC Solution is critical for scientific research, governmental affairs, industry, and national defense. It is often the sole or major solution for costly, complex theoretical models or experiments. Leveraging its strong capability to innovate and product design chip, hardware and system architecture, Huawei develops a variety of advanced IT infrastructures such as servers, storage and networks with faster performance, higher energy-efficiency and more flexible management.

Test configuration details

In this document, we report the performance test results of ANSYS Fluent release 17.2 on Huawei HPC solution.

ANSYS Fluent benchmark cases

Among the several benchmark cases applicable for simulation scenarios in various industries, sedan_4m and ice_2m are two typical benchmark simulation scenarios in the automotive industry. They simulate and analyze external flow over a passenger sedan and in a 4-stroke spray guided gasoline direct injection model.




Figure 1: ANSYS Fluent Benchmark case 1: sedan_4m

Number of cells	4,000,000
Cell type	Mixed
Models	Standard K-e Turbulence
Solver	Pressure based coupled solver, Green-Gauss cell based, steady

Table 1: Configuration of sedan_4m




Figure 2: ANSYS Fluent Benchmark case 2: ice_2m

Number of cells	2,000,000	
Cell type	Mixed	
Models	Standard K-e turbulence, DPM, species partially premixed combustion Dynamic mesh	
Solver	Pressure-based coupled solver, Green-Gauss cell-based, unsteady	

Table 2: Configuration of ice_2m

Details of the system configuration are described in Table 3

Huawei HPC Cluster Configuration

Table.3 Configuration of HPC

SYSTEM UNDER TEST	HUAWEI FusionServer X6800 High-density Server (Including 16 pcs XH620 v3 blade nodes)
PROCESSORS	2x Intel Xeon E5-2680 V4 (14 core, 2.4GHz)
MEMORY	8x 16GB DDR4-2400MHz
NETWORK	56G FDR IB network
HARD DRIVE	2x 300G 10K RPM
OS	Red Hat Enterprise Linux 6.7
MPI	Intel MPI 5.0.3, Open MPI 1.6.5, IBM Platform MPI 9.1.3.1

Featuring high density and simplified system management, Huawei X6800 offers flexible configuration of a wide range of server nodes, meeting differentiated requirements for computing, storage, and I/O. The server for this test was configured with the maximum of eight XH620 blade nodes per frame, supporting Intel Xeon E5-2600 v3 and v4 series processors, 16×DDR4 DIMM slots, 2×Disk drives and FDR/EDR interconnection.



Figure 3: Huawei X6800 High-Density Server

Test Results

The performance test results are evaluated in Table 4

Table 4: Quotas

Quotas	Descriptions
Solver Rating	<p>A measure of the number of tasks that can be done in 24 hours. Higher scores indicate better performance. The solver rating is calculated as follows:</p> <p>Total number of seconds in a day (86,400) divided by the time of the solve step in the job.</p> <p>For example, if the solver time in a particular job run were 100 seconds, the solver rating for that job would be 86,400/100, which is a solver rate of 864.</p>
Speedup	The ratio of wall-clock time required to complete a given calculation using a single processor to that of the equivalent calculation

The following tests were conducted:

- Performance on a single node
 - Scalability
 - With turbo On and Off
 - With memory frequency (2400/2133/1867 Hz)
- Performance on multi-nodes
 - Scalability
 - With GE and FDR
 - With IntelMPI, OpenMPI and PCMPI

Performance on a single node

The tests were conducted to evaluate the solver rating and speedup from 1 core to 28 cores.

Turbo mode: makes it possible to automatically over clock the cores under certain conditions. HT: Hyper-Threading is a technology that allows a single microprocessor to act like two separate processors in the operating system and the application programs that use it. With Hyper-Threading, a microprocessor's core processor can execute two (rather than one) concurrent streams (or threads) of instructions sent by the operating system.




Figure 4: Solver ratings and speedup on a single node (sedan_4m)




Figure 5: Solver ratings and speedup on a single node (ice_2m)

Figure 4 and Figure 5 show that ANSYS Fluent version 17.2 has good scalability when increasing the quantity of cores in a single node. The speedup of case Sedan_4m is nearly linear.




Figure 6: Solver rating with Turbo on and off (sedan_4m)




Figure 7: Solver rating with Turbo on and off (ice_2m)

Figure 6 and Figure 7 show that Turbo mode affects the performance in these cases, especially in the ice_2m case, meaning higher frequency leverages better performance.

With Memory frequency 2400/2133/1867




Figure 8: Solver ratings with different memory frequency

As Figure 8 illustrates, the memory frequency is a critical factor influencing the performance when running these benchmark cases on any amount of cores. The higher the memory frequency, the better the performance.

Performance on multi-nodes

Scalability




Figure 9: Solver ratings and speedup on multi-nodes (sedan_4m)




Figure 10: Solver ratings and speedup on multi-nodes (ice_2m)

As Figure 9 and Figure 10 illustrate, the scalability of clusters are excellent from 1 node (28 cores) to 16 nodes (448 cores). The results show sedan_4m has good scalability with increasing amount of nodes. Speedup is nearly linear. However, ice_2m grows gently along with the nodes, especially for more than 4 nodes (112 cores).

With GE and FDR

The network is critical in determining the scalability from the hardware side. With parallel process applications, there are substantial amounts of communication going on between each process in a node and also with processes spread out over multi-nodes in the network. Because of this, a network with low latency and high bandwidth is required for performance.




Figure 11: Solver ratings on Infiniband (FDR) and Ethernet (GE) (sedan_4m)




Figure 12: Solver ratings on Infiniband (FDR) and Ethernet (GE) (ice_2m)

As Figure 11 and Figure 12 illustrate, small cluster systems (four nodes or less) connected with a GE network may provide sufficient performance. For large cluster systems (more than four nodes), Infiniband (FDR) is the better interconnect method, providing a high-throughput, low-latency fabric for node-to-node communications.

With IntelMPI, OpenMPI and PCMPI




Figure 13: Solver ratings with different MPI (sedan_4m)




Figure 14: Solver ratings with different MPI (ice_2m)

(Version: Intel MPI 5.0.3, Open MPI 1.6.5, PC MPI 9.1.3.1)
The performance of three MPI types is much the same when running these benchmark cases on 1 node and 2 nodes. PC MPI has the best performance when running these benchmark cases on 8 nodes and 16 nodes.

Conclusions

The above tests indicate that ANSYS Fluent with Huawei X6800 server has optimal scalability within a single node or over multi-nodes in an HPC cluster.

The quantity of cores and frequency of processor have significant influence on the applications performance. It is necessary to consider when selecting a processor model that the license model and processor specifications. Because high frequency results in high performance, the memory frequency could be the determining factor in improving performance.

When expanding the multi-nodes, the task is dispatched to each node. The delay in interconnection and bandwidth affects the performance. Different MPI protocols influence the performance as well.

The current tests were focused on helping customers achieve the configuration with the best performance-to-price ratio when choosing Huawei HPC solution with ANSYS software for different application and scenarios. In the future, Huawei will keep testing more ANSYS solutions in several scenarios. In order to offer the best HPC solution for automotive simulations, Huawei will continue to optimize the Huawei HPC solution based on automotive application requirements.