Best Practices for Implementing ANSYS® Fluent® Software on Cluster Technologies from IBM

Highlights:

- Optimize performance of ANSYS Fluent software on cluster technologies
- Select the right cluster components for small, medium and large implementations

ANSYS® Fluent® is a computational fluid dynamics (CFD) software solution used to predict fluid flow, heat and mass transfer, chemical reactions, and related phenomena by numerically solving a set of governing mathematical equations (conservation of mass, momentum, energy and others). ANSYS Fluent, along with other engineering simulation tools from ANSYS, helps engineering teams understand product performance during conceptual studies of new designs, product development, troubleshooting and redesign.

Scaling up high-performance computing resources allows engineering teams to dramatically increase the value of engineering simulation, enabling more detailed and accurate simulations and more design exploration within the time constraints of project deadlines. Choosing the right HPC resources, however, requires an understanding of how key technologies determine application performance. In this short document, ANSYS and IBM share conclusions about how to optimize the performance of ANSYS Fluent software on cluster technologies.

Please visit http://ansys.com/About+ANSYS/Partner+Programs/HPC+Partners/IBM for more details and recommended configurations for engineering simulation with ANSYS Fluent.

Processor Clock Speed

There is always an improvement (5% to 10%) in application performance when clock speed increases. Compute nodes based on slightly slower clock are recommended because the actual performance improvement when using the fastest available clock is very small. While there is no harm in choosing the fastest processors, investing in other aspects of the cluster is likely to be more cost-effective.

CPU Turbo Boost

The Turbo Boost should be left on to extract more performance from the processors.

Hyper-Threading

Hyper-Threading can improve performance, but only by a relatively small amount (about 3% to 8%). It will also consume twice as many licenses. If ANSYS Fluent customers have a limited number of licenses, Hyper- Threading is not recommended

Sockets

A 4-socket system is recommended if the performance requirements of a single ANSYS Fluent job can be satisfied with a single 4-socket system. If capacity needs to scale beyond a 4-socket system, it is more cost-effective to buy 2-socket servers in order to build a cluster with high-speed interconnect, such as 10-gigabit Ethernet or Infiniband.

Processor Core Density

If the primary consideration is optimizing performance for a fixed number of ANSYS Fluent licenses, systems with 4-core processors are recommended because they provide better performance — typically 20% to 30% over systems using 6-core processors. However, if the primary consideration is total cluster cost, 6-core systems are a better choice for an equivalent total core count because they cost significantly less than 4-core systems.

Memory Configuration

For ANSYS Fluent, variations in performance due to memory speeds can be significant. To enable memory to operate at the maximum possible speed, populate all channels with an equal amount of memory. Usually, 24 GB memory per node is sufficient because ANSYS Fluent workloads requiring more memory would typically be run on multiple nodes.

Number of Nodes

For high-end clusters, performance levels off while total cost rises linearly with the size of the cluster. Low-end clusters may not satisfy productivity requirements. For a given problem size, the optimal number of nodes in the cluster typically falls in the middle of the performance/price spectrum. The choices within this narrow range should be evaluated closely (using ROI analysis) to determine optimal cluster size.

Interconnect Selection

For cluster sizes up to 8 nodes running simulations that are a few million cells in size, a 10-Gigabit network is recommended. When considering larger clusters, Infiniband will optimize performance and offset additional costs by reducing the number of nodes required for a given level of performance.

Infiniband

For most Infiniband networks, use QDR Infiniband with 4X binding. Because ANSYS Fluent is sensitive to network latency, it will benefit from the low latency of Infiniband networks. ANSYS Fluent does not seem to be very sensitive to the blocking design of Infiniband networks. A blocking factor of 2:1 or even 4:1 offers an ideal balance of application performance and total investment.

© Copyright IBM Corporation 2012

IBM Global Services
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
April 2012
All Rights Reserved

IBM, the IBM logo, ibm.com and System x are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: **ibm.com**/legal/copytrade.shtml

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other product, company or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

Please Recycle

For more information

ANSYS
hpcinfo@ansys.com
www.ansys.com

IBM HPC
loribron@ca.ibm.com
http://www-03.ibm.com/systems/deepcomputing

XSO03161-USEN-00