

## General Hardware Recommendations

We test and certify our applications on high-end workstations or servers and partner with all of the major hardware vendors. Below is the current list of certified recommended processors. At the hardware vendor web site, look for " **Workstations** " or " **Servers** " when looking for these processors. We encourage you to configure a system and send us the specifications and we will look them over and get back to you with further recommendations, if necessary.

### VENDOR SOLUTIONS

<http://www.ansys.com/About%20ANSYS/Partner%20Programs/Complete%20Windows%20Support%20for%20ANSYS>

### PROCESSORS

Intel four - six-core (Westmere, X56xx) Processors  
AMD quad-core Shanghai processors

**NOTE:** The absolute best scaling you can get is with fast clock speed

### BENCHMARKS

<http://ansys.com/Support/Platform+Support/Benchmarks+Overview>

### ANSYS Platform Support

<http://www.ansys.com/Support/Platform+Support>

[ANSYS Platform Support Strategy \(PDF\)](#)

### HYPER THREADING (Simultaneous multithreading or SMT)

Hyper-Threading is now called Simultaneous multithreading or SMT. Customers are recommended to leave SMT enabled on their systems but not over-subscribe physical cores for parallel simulations. While some improvement is possible, the extra performance from the virtual threads is not cost-effective and incommensurate with the additional license costs (which are per process)."

Basically, if a section of the CPU core is not being used it tries to run a second task on these sections. For example, if one process only needs to do floating point operations while another only needs to do integer operations they can run both concurrently. For FLUENT, there is no consequence to performance if it is turned off. If SMT is on, and you run 16x (instead of 8x; assuming dual cpu quad-core nodes), you can get an additional 20% or so (compared to 8x) improvement. This is not recommended since you only get 20% more for 2x licenses (license is per process). in this scenerio, leave SMT on and run 8 way. This is the recommended approach.

**ANSYS Mechanical:** Mechanical does not benefit much from SMT enabled and it is not worth purchasing the extra HPC licenses for the small gain that might be realized.

### GPU COMPUTING

The use of graphics processing units, GPUs, to perform general purpose computations is becoming more prevalent. GPUs offer a highly parallel design which often includes hundreds of cores (sometimes referred to as shaders) and have their own dedicated physical memory. Certain high-end graphics cards, the ones with the most amount of cores and memory and the ones that can perform double precision computations, can be used to accelerate the computations performed during an ANSYS simulation.

It is important to note that not all graphics cards meet these requirements. For details on which GPU devices are supported and the corresponding driver versions, see the GPU requirements outlined in the Windows Installation Guide and the UNIX/Linux Installation Guide.

### ANSYS Mechanical

ANSYS Mechanical 13 can use a single GPU (only) when the SMP solver is invoked. That is, the GPU is

**not** enabled for the "D-ANSYS" distributed/cluster solver.

## GPU REQUIREMENTS

Your system must meet the following requirements to use the GPU capability in Mechanical APDL.

1. The machine being used for the simulation must contain at least one **nVIDIA Tesla series GPU card** (a Tesla 20-series card is recommended for optimal performance).
2. On Windows, the driver version for the nVIDIA Tesla series GPU card must be 258.96 or newer.
3. You must be running on a Windows 64-bit or Linux x64 operating system.
4. Before activating the GPU accelerator capability, you **must** have a supported GPU card and at least one **ANSYS HPC Pack license**.

**NOTE:** On Windows, the use of Remote Desktop disables the use of a GPU, which means the GPU Accelerator Capability cannot function under Remote Desktop. Alternative means of remotely logging into a machine can be used (e.g., third-party software such as VNC or other).

**NOTE ABOUT TESLA AND VISUALIZATION:** It is necessary to have another nVIDIA display card in the system for visualization.

For information on activating the GPU please see ANSYS Help, specifically:

### 16.1. Activating the GPU Accelerator Capability Open the Mechanical APDL Product Launcher.

## ANSYS CFD Solvers

ANSYS CFD solvers do not currently run on GPU processors.

For more information on **nVIDIA** Tesla visit: [http://www.nvidia.com/object/tesla\\_computing\\_solutions.html](http://www.nvidia.com/object/tesla_computing_solutions.html)

## HARD DRIVE SPECIFICATIONS

**FLUENT/CFX** - No special hard drive configuration is required as FLUENT does very little Disk Read/Write.

**MECHANICAL** - Multiple SCSI or SAS hard drives striped with RAID 0, 15,000 RPM or better is recommended. To clarify, if you have enough memory for the sparse solver to solve in-core, which is quite possible these days, or you use the PCG solver, then the I/O is essentially the same as FLUENTs.

### RAID 0 Description

Using RAID 0 mainly for tasks requiring fast access to a large capacity of temporary disk storage. RAID 0 usable capacity is 100% as all available drives are used.

### MEMORY

We recommend a minimum of 4 - 8 GB RAM per core.

## HIGH SPEED INTERCONNECTS

In general if you have fewer than 4 nodes, a Gigabit Ethernet switch is sufficient. For more than 4 nodes, consider a high speed switch such as Infiniband or Myrinet. The amount of speedup you see from such a switch will depend on your problems size and complexity. Currently, Infiniband and Myrinet are only available on Linux systems or Windows 2008 HPC Server (High Performance Computing).

**Infiniband Notes about FLUENT:** In general, if the cluster is relatively small say 16 nodes or less and was running just FLUENT then DDR InfiniBand would be adequate.

## SUPPORTED INTERCONNECTS

Please visit this link for a list of supported interconnects:

<http://www.ansys.com/Support/Platform+Support/ANSYS+13.0+Supported+Interconnects>

## GRAPHIC CARDS

We test and certify with the following two graphic card vendors. These cards range from the low end to the very high end. We recommend the high end

nVIDIA Quadro and Quadro FX  
ATI FireGL

**NOTE ABOUT TESLA AND VISUALIZATION:** If you are using a Tesla C2050 or C2070, it's not necessary to have another nVIDIA display card in the system, but it is highly recommended. If you are using a Tesla M2050 or M2070, you need to have another nVIDIA card for visualization as the Mxxx series does not have video out.

## REMOTE ACCESS - BEST PRACTICES

Running any ANSYS application using remote access software is **not certified or supported**. If the machine you are remotng into has a PCIx16 slot with an nVidia Quadro FX or AMD ATI FirePro/FireGL high-end graphics card it could work with some considerations. Remember that most servers **DO NOT** have a PCIx16 slot so this will not be possible. There are servers on the market that do offer this capability. Please check your hardware vendors web site for more information.

Visit this FAQ for recommendations and best practices.

<http://www.fluentusers.com/support/installation/winfaq/remote-access.htm>

## SUBMITTING JOBS TO A CLUSTER

### WINDOWS > WINDOWS CLUSTER

Having a high-end 64-bit Workstation to do Pre and Post Processing locally is supported. Pre and Post Processing has heavy graphics demands and requires a high-end, certified video card. See the section above, "**REMOTE ACCESS - BEST PRACTICES**" and choose the FAQ link for more information.

### LINUX > LINUX CLUSTER

Having a high-end 64-bit Workstation to do Pre and Post Processing locally is supported. Pre and Post Processing has heavy graphics demands and requires a high-end, certified video card. See the section above, "**REMOTE ACCESS - BEST PRACTICES**" and choose the FAQ link for more information.

## 64-BIT OPERATING SYSTEMS

We recommend that you run on 64-bit Workstations or Servers. There is a 2 GB per process memory limitation on 32-bit operating systems. Realistically, approximately 1.5 - 1.7 GB of RAM is all that will be available. Running on 64-bit processors will give you access to as much RAM as is configured on the system. To see information about Microsoft Windows Operating Systems memory limitations visit: <http://msdn.microsoft.com/en-us/library/aa366778.aspx>

[Home](#)