Recommended hardware system configurations for ANSYS users

The purpose of this document is to recommend system configurations that will deliver high performance for ANSYS users across the entire range of ANSYS user experiences. In recent years, the performance of processors has become nearly identical at every price point from inexpensive desktop systems to expensive multi-processor servers. Yet there can still be significant performance bottlenecks that add hours or days of execution time to ANSYS runs. The main cause of the problem is rarely slow processors; rather, it is a severe system imbalance from either memory size limitations, poor I/O setups, or slow interconnects between processors on distributed memory systems. In the world of commodity-priced computing systems, salespersons often wear many hats and may lack the experience or knowledge of the demands of a compute intensive application like ANSYS. ANSYS analyses involve ever larger and detailed models coupled using complex and sophisticated numerical simulations that fully utilize system CPU, memory, and I/O resources.

Users who set their modeling expectations within the limits of their available computing resources can benefit from maximum performance at every price point. At the same time, each system configuration has limits beyond which the hardware performs well below its maximum performance expectations. The following criteria are a guide for users who select systems at each level of hardware price and configuration. The compute demands for various model size estimates can vary greatly depending on the analysis type, solver choice, and element selection.

First, some general memory guidelines for ANSYS.

1) The most general guideline for ANSYS solver memory is **1 GB per million degrees of freedom**. This is true for the iterative (PCG) solvers as well as the direct sparse solver. Solver memory is the high water memory usage for ANSYS so the 1 GB per million DOFs is a good starting estimate for total ANSYS memory usage for most runs. Blocky 3-D models using higher order elements require more memory (2x to 3X more in the case of the sparse solver) and a few commonly used elements in ANSYS (SOLID92 and SOLID95 for example) may use significantly less memory using the PCG iterative solver due to the MSAVE,ON option (now the default setting, where applicable)

2) I/O requirements are the same as memory requirements for iterative solver jobs except that some analyses with multiple load steps may generate very large results files. Sparse solver runs require **10 GB per million DOFs for files**. This number may also

grow by 2x to 3X for dense 3-D models using higher order elements. Block Lanczos analyses are among the most demanding runs in ANSYS, requiring a combination of CPU and I/O operations.

3) I/O cannot keep up with CPUs. This is a manageable problem with iterative solvers but it can result in severe performance degradation for sparse solver runs. Most modern operating systems, including Windows, now automatically use extra memory for a large file cache. **Large memory systems are therefore advantageous at every price point.**

4) High performance solution on any hardware configuration can be obtained when the ANSYS model size is such that the ANSYS memory requirement is comfortably satisfied within the available system memory.

5) ANSYS has demonstrated the capacity to utilize 8 to 16 processors on systems exceeding 100 GB of memory. On such systems, ANSYS can solve 100 million DOF size models, leading our industry both in capacity and time to solution. On smaller systems, ANSYS has intentionally designed memory management for high performance so that it is possible to run ANSYS solutions at near-peak performance on properly configured systems, subject to the guidelines above on problem size relative to available memory.

Hardware recommendations to maximize ANSYS performance:

1) **Desktop 32-bit systems** (primarily Windows 32-bit OS and some Linux 32-bit)

Any desktop system for ANSYS should have a **minimum of 2 GB of real memory. Swap space should be set for 3 GB** and users should configure Windows XP systems to boot with the **/3GB switch enabled**. This configuration should be done in a way that allows users to toggle this capability because there are some known problems in some cases. Disk resources should be **at least 100 GB of space**. This space should be on a separate drive used as a scratch space when running large models, with permanent files moved to another drive or offline after the run completes. This drive should be cleaned out and even reformatted regularly to limit the amount of file corruption and fragmentation that is famous on Windows systems.

10,000 RPM or 15,000 RPM drives should be used. Currently there are few systems at this price point that deliver sustained I/O performance greater then 50 – 70 MB per second for large ANSYS jobs. The best thing to do is avoid models that would do excessive I/O. For the PCG solver, these systems should be able to solve 1 million to 3 million DOFs efficiently with high performance. For sparse solver runs, the best performance comes running incore but this mode will be restricted to 250k DOFs and less in most cases. It is possible to run 1 million DOFs with the sparse solver but there will be significant I/O and for blocky 3-D models performance will be very degraded.

The **32-bit desktop system is rapidly being replaced by AMD Opteron and Intel em64t systems**. Both systems run 32-bit applications as well as 64-bit applications. Windows is now released for 64-bit systems and the 2 GB limitation of the Windows32 environment need not constrain even commodity systems anymore. There may be some delay in fully implementing Windows64, depending on the speed of third party software and hardware suppliers in producing drivers and other software for Windows64. At any rate, one can purchase 64-bit commodity priced systems and they will be upgradeable to full 64-bit capability with software updating.

2) Desktop 64-bit systems (Opteron, Intel em64t, and Intel Itanium workstations)

These systems should have a **minimum of 8 GB of memory and 2 processors**. Disks should be **10k or 15K RPM drives with at least 200 GB** of free space. The ANSYS jobs should run in a separate, striped partition. 8 GB will allow sparse solver jobs of 500k to 750 k to run incore or very efficiently out-of-core. Most sparse solver jobs, even up to 1 million or 2 million DOFs will be able to run in an efficient out-of-core implementation that will not degrade performance like 32-bit OS experience. This is a sweet spot for these systems over 32-bit machines. An additional benefit of 8 GB of memory for smaller systems is much better elapsed time performance for those jobs that nearly fill memory on a 32-bit machine. Even jobs with 300-500 k DOFs may run entirely incore, allowing the system to run at sustained peak performance.

Most of these systems still lag in delivering good sustained I/O performance. An exception is the HP Itanium workstations which could be configured with 3 striped SCSI drives for over 200 GB of space. These systems sustain over 100 MB per second in sparse solver I/O intensive solves. Unfortunately, HP is no longer producing these Itanium-based desktop workstations. It remains to be seen if em64t systems or Opteron systems will be configured to match or exceed this I/O performance. In the meantime it is **best to run jobs whose memory usage is well within the 8 GB limit, thus avoiding the most costly I/O.**

These systems can use 2 processors effectively but they generally **should be considered single job/single user resources**. They generally do not have the OS refinements nor resources to support multiple users or more than one demanding job at a time.

3) Deskside Linux and Unix systems

These systems can **extend the desktop 64-bit performance to 4 or even 8 processors** and extend available memory well beyond 8 GB. They can support more than one user or multiple jobs, as long as the user is careful to not oversubscribe processors or memory. For multiple jobs and users, it is even more important to leave some memory free to function as a buffer for I/O. **A 4-CPU system should have at least 16 GB of memory, preferably 32 GB**. It should have a **fast striped array of 5 or 6 disks with a capacity of 250 to 500 GB**. A system like this can do multiple routine ANSYS jobs at peak speed, service more that one user, and also solve occasional “hero-sized” jobs with several million DOFs. The iterative solver could easily handle 20 million DOFs on such a system and the sparse solver could handle 2 – to 5 million DOFs running at sustained peak performance. Distributed ANSYS will efficiently use all of the CPUs for fastest elapsed time and will run very well due to the speed of shared-memory servers on MPI applications. Additionally, the large shared-memory image of these systems will make pre- and post-processing of very large models tolerable. Users should allocate large portions of memory to database space for post processing of very large models to eliminate the need for page file accesses during graphics operations.

Rack mounted servers with large memory and fast I/O are some of the most cost effective systems available today using opteron and em64t hardware and Linux OS. They seem to be better balanced with larger memory and much better I/O than the current desktop systems that use the same hardware. The price difference is recovered by a nearly equal differentiation in solution time on the large problems.

4) High end servers

These systems offer maximum performance all the way up to the current limitation of ANSYS software. They should be configured with maximum memory per processor. It is better to buy fewer processors and more memory on such a system. An 8-processor system with 64 GB of memory is a better system to run ANSYS than 16 processors with 32 GB of memory. An ideal configuration for a high end server system is 16 processors (Itel Itanium, 1.6 Ghz) with 128 GB of memory and 1 TB of striped disks. This system can solve multiple million DOF-sized jobs, a single 100 Million DOF analysis or a steady workload of ANSYS jobs including remote solves from ANSYS Workbench. It runs shared memory parallel and achieves even better parallel performance using PPFA and distributed MPI-based parallel performance. This is the only system configuration that can solve a 100 Million DOF sized model currently and it is also the only system configuration that can support multiple very large jobs in excess of 1 million DOFs all running at peak performance.

5) Cluster Systems and Distributed ANSYS specifics

Cluster systems are rapidly evolving and deserve separate comments. Just as ANSYS requires high performance I/O, the distributed parallel ANSYS code, Distributed ANSYS, **requires distributed systems with high performance interconnects**. The amount of data passed between processors in a typical ANSYS run is measured in hundreds of megabytes to gigabytes of data with the number of messages often exceeding several hundred thousand messages.

All of the hardware options listed above in 1 – 4 can be configured for Distributed ANSYS with the exception of Windows systems. ANSYS is still working to find a suitable MPI interconnect with fast MPI libraries, which allow the power of parallel processing on multi-CPU and multiple machine Windows clusters to be realized. It is expected that Windows64 will produce a competitive MPI library for clusters of Opteron or em64t processors.

Individual ANSYS jobs using Distributed ANSYS rarely see significant scaling past 8 processors. As in classic ANSYS runs, **it is better to have powerful nodes with plenty of memory than to have lots of cheap processors with limited memory**. The **host node on these systems should have 8 GB**. **All other nodes should have at least 4 GB**. Clusters of 32-bit processors dominate the market today but for the additional memory demands of ANSYS, 64-bit clusters will be much more effective. It is possible to extend the capacity of 32-bit systems if fast interconnects are used. At a minimum one needs **Gige type connection for Distributed ANSYS runs with faster interconnects like Miranet and Infiniband preferred**. Each ANSYS process will do I/O in addition to extensive message passing. Therefore **each processor should have local disks or a separate network connection that allows very high speed I/O** transfer back to a large central I/O system. This configuration is typically seen on higher end large scale cluster systems. The individual disks do not have to be huge. **15 – 20 GB should be more than adequate in most cases**. However, the **host node needs large I/O capacity** because all results come back to the host node in Distributed ANSYS runs and initial input files are stored and read from the host node.

Poor interconnects between true cluster systems will result in solution time going up as processors are added. Large shared memory servers have the best MPI performance in many cases. The large shared memory image is advantageous for pre- and post-processing and also accommodates the large volume of message passing in ANSYS, all at memory copy speeds rather than much slower network connection speeds. Good Miranet and Infiniband interconnects can approach shared memory interconnect speeds and other fast interconnect technologies are also emerging.

In summary, Distributed ANSYS system recommendations are: Follow the same basic recommendations given in 1-4 for ANSYS systems. Additionally, configure adequate local disks for each processing node and use high performance interconnects. Remember that regular internet network connections are not sufficient for Distributed ANSYS speedups. Linux and UNIX systems are the only useful options currently but ANSYS is working with Microsoft to find a competitive solution for Windows environments.