

**If you're new to high-performance computing or parallel processing, the following nomenclature may be helpful:**



- **HPC:** HPC stands for High Performance Computing and refers to the use of today's multiprocessor computers to solve the most challenging computational problems in science and engineering.
- **Parallel Processing:** Any technique that divides a problem into pieces and simultaneously solves each piece using multiple computers or multiple processors. Parallel processing can also refer to using multiple computers or processors to simultaneously solve a group of independent problems.
- **Multi-threading:** A parallel processing technique that splits data and tasks into parallel subtasks and let the underlying architecture manage how the threads run, either concurrently on a single core or in parallel on multiple cores. Multi-threading for geometry and meshing operations does NOT require any HPC licenses!
- **Domain Decomposition:** This is a particular technique for parallel processing, in which the simulation model is divided (decomposed) into pieces and each piece of the model is sent off to be computed on a separate processor. Many (most) of the ANSYS solvers use this domain decomposition method.
- **Machine / Socket / Processor / Core:** This is the hierarchy of computer technology. A "machine" might be a workstation or a single server (sometimes called a "node" or "blade") in a cluster. Each machine will have one or more "sockets". Each socket contains a processor (also called a CPU). Each processor typically contains multiple cores. It's the cores that ultimately do the processing. ANSYS HPC licensing controls the number of cores that can be used.
- **Workstation:** Graphics-enabled desktop or deskside computer. Today, workstations may include up to 16 cores (or more) with lots of memory (64GB RAM or more), and can be considered "HPC" for a single user. Workstation users can benefit from ANSYS HPC!
- **Compute Cluster:** A compute cluster is a group of server units (also called nodes or blades) that are connected together (by interconnect technologies such as "gigE" or "InfiniBand") and can work together on a single task. This is generally considered the best scalable hardware solution to enable a group of users or multiple simultaneous simulations.
- **Schedulers:** These are mechanisms to allow single or multiple users to submit tasks to a computer or set of computers without having to worry about the explicit details of how or where the job will run. Popular schedulers include Platform LSF, PBS, Sun Grid Engine (now Oracle Grid Engine), and Microsoft HPC 2008 (previously Microsoft CCS).
- **RSM:** The Remote Solve Manager is used primarily for long-running simulations that do not fit within your workstation's resources. Through submission to RSM, the solution can be executed on remote computing resources. RSM can also submit jobs to the local machine to allow the queuing of solutions on your workstation.

- GP-GPU (or GPU):** General Programmable Graphical Processing Units, or GP-GPUs (GPUs, for short), are an extension of the processors used on graphics boards, but with a new focus on computational throughput. GPUs have the potential to be a game changer, with 100's of low-power cores on a single processor and the ability to perform an order of magnitude more operations per second than current multicore CPUs. In general, throughput for ANSYS applications on GPUs is a technical challenge, primarily due to bottlenecks related to memory access. We are engaged with NVIDIA, AMD (ATI), and Intel on GPU R&D. ANSYS introduced GPU capability to accelerate ANSYS Mechanical simulations using the shared memory solver and GPU innovations (like support for distributed memory solvers and multiple GPUs within a cluster). GPUs can be used to accelerate radiation heat transfer calculations.
- DSO:** DSO (Distributed Solve Option) is a HPC product offering for our Electronics products only. DSO provides licensing that enables multiple simulations to run simultaneously, in order to examine the effect of parameter variation, geometry variation, or frequency variation. This kind of “multiple design point” product packaging (pricing) is also available for our Structures/Fluids products in the form of HPC Parametric Packs.
- RDO:** Robust Design and Optimization (RDO) refers to the use of multiple simulations to explore the impact of design uncertainties (robust design) or to arrive at an improved design that achieves certain objectives (optimization). RDO is a HPC topic because RDO requires a large number of independent (but linked) simulations that can only be performed with sufficient HPC resources.
- Cloud Computing:** Cloud computing refers to using hardware and/or software that is accessed over the internet and typically rented (vs owned) or paid for on a usage basis. ANSYS customers may be asking about our “cloud” strategy. A simple answer is that we can support customers who want to use remote HPC hardware resources (e.g., accessed in the cloud). Customers can use their “normal” licenses to run on remote hardware and/or they can buy short term licenses for this purpose. You may need to consider LAN/WAN policy issues, in order to set this up, but we want to reassure customers that we can support their use of HPC resources wherever they can access them. For more details, please visit our Cloud & IT Solutions section.

