White Paper

The Value of High-Performance Computing for Simulation

High-performance computing (HPC) is an enormous part of the present and future of engineering simulation. HPC allows best-in-class companies to gain high-fidelity insight into product behaviour, insight that cannot be obtained without the detailed simulation models – including more geometric detail, larger systems, and more complex physics. When applied to design exploration that enables a product development organization to explore how a design will behave across a range of real-world operating conditions, HPC can lead to robust product performance, and reduced warranty and maintenance costs.

This value only can be delivered by a continuous software development focus on HPC. Companies invest heavily to ensure that their products deliver truly high-performance results on the latest HPC architectures - possibly using hardware accelerators like NVIDIA GPUs and Intel Xeon Phi co-processors [1] – and at an ever-increasing scale of parallelism being delivered through shared-memory parallel (SMP), distributed-memory parallel (DMP), hybrid SMP/DMP, and "embarrassingly parallel computing" capabilities.

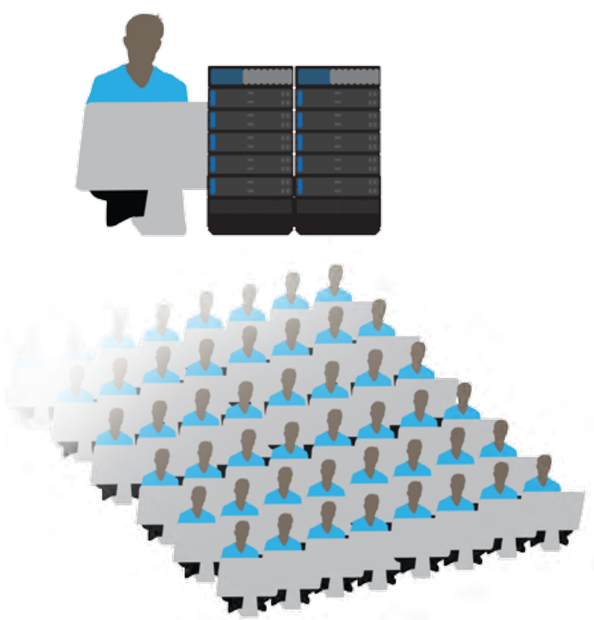

Figure 1: HPC capability and HPC capacity reflected by two different HPC usage scenarios: single user running a large job on a cluster (top image) and many users running smaller jobs on a workstation (bottom image).

HPC Capability vs. HPC Capacity

The ability to run a single simulation job on 2,000 computing cores clearly delivers a different return on investment (ROI) than accommodating 100 engineers each running on 20 computing cores, as schematically shown in Figure 1. The first scenario is referred to as HPC capability while the second is often called HPC capacity. These two different HPC usage scenarios require different HPC pricing and packaging models.

HPC Packs – Scalable Pricing for Capability Computing

Capability computing applies a large number of cores to execution of a single simulation job. ANSYS software is architected to show near-ideal (i.e. linearly increasing) solution speedup[1] with increasing number of cores. With every doubling of core count, the time to solution, ideally, is cut in half. While the core count is doubled over and over, however, cores begin to cut minutes in half instead of hours or days. The key is to have enough cores at your disposal to solve the detailed problem you need to solve within the time you require, all while recognizing that there is a diminishing value of added cores after some point.

[1]While ANSYS development is focused on increasing the linear region of scalability, the actual speedup may depend on various factors including hardware configuration, simulation model size, and model physics involved.

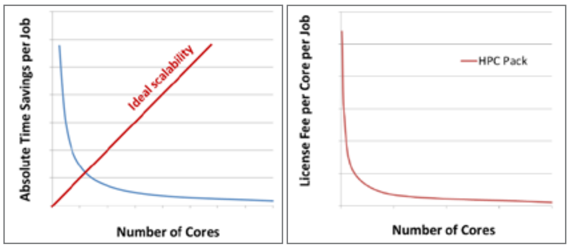

Figure 2: Even with ideal linear scaling, time to solution changes as more cores are added (left image). HPC Packs reflect this value proposition by driving the incremental cost of more cores to near-zero at high core counts (right image).

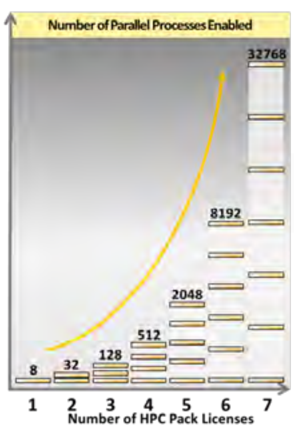

Figure 3: HPC Packs deliver virtually unlimited parallel capability for a fixed incremental cost.

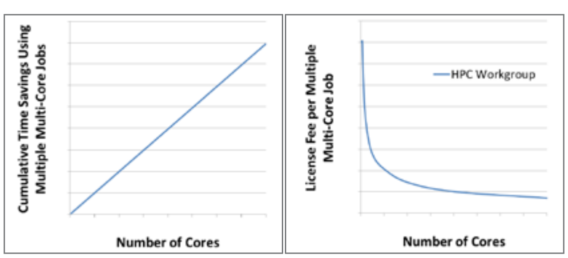

Figure 4: Computing capacity example where HPC applied to multiple multi-core jobs creates cumulative value that linearly increases with core count (left image). The HPC Workgroup licensing model rewards volume buyer with scaled pricing
(right image).

Consider an example. If a CFD simulation takes 16 days on a single core, adding 32 cores might reduce the simulation time to four hours, slashing out more than two weeks. Doubling the number of cores from 32 to 64 might cut the time in half again – from four hours to two hours. Clearly the incremental value of parallel decreases as the number of cores increases, and a value-based pricing model needs to accommodate this idea.

ANSYS HPC Pack and ANSYS Electronics HPC Pack are the HPC licensing solution from ANSYS for capability computing.

HPC Packs provide:

- ***Virtually unlimited HPC at a fixed incremental cost:*** You decide what unlimited means to you and buy the right number of HPC Packs (see Figure 3) to support whatever level of capability computing you need. HPC Packs enable highly scaled parallel for the most computationally demanding simulations.
- ***Scalable cost model:*** In keeping with the law of diminishing returns noted above, you pay less as you add more capability.

Running on 2,000 cores instead of 20 cores incurs a cost premium of 1.5 times and not the 100 times that the core increase might imply.

HPC Workgroups – Volume Pricing for Capacity Computing

If the value of 1000s of cores applied to a single job asymptotes as shown in Figure 2, what is the value of those 1000s of cores applied to multiple jobs and multiple users? The answer is depicted in Figure 4, which shows productivity - measured in cumulative time saved - as multiple jobs are executed using for this kind of capacity computing (also called embarrassingly parallel computing), the ROI increases in a linear fashion with the number of cores used because the cores enable multiple jobs to execute with faster turnaround. This faster turnaround creates increased simulation throughput, increased design understanding and reduced time to market while maintaining reliable product performance.

When HPC is used for capacity — the ability to run multiple jobs - the ROI, accordingly, increases in a linear fashion with the cores.
ANSYS HPC Workgroup and ANSYS Electronics HPC Workgroup provide volume access to parallel processing for increased simulation throughput for multiple jobs and/or multiple users:

- The pool of HPC licenses is sharable between multiple users (or jobs) with increasing ROI as more cores are used to create more simulation throughput.
- Volume pricing rewards the volume buyer. For example, filling up a 32,768- instead of a 2,048-core cluster with 128-core jobs will cut the price per job in one quarter.

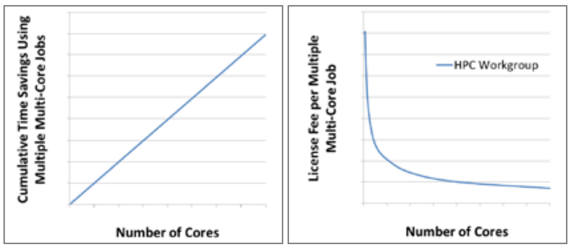

Figure 5: Pricing model of HPC Parametric Pack reflects a decreasing cost per design point as the number of design points in a single study increases.

Figure 6: Similar to ANSYS HPC Packs that deliver scalable parallel processing, HPC Parametric Packs scale functionality when the customer chooses to use more than one HPC Parametric Pack per design study.

HPC Parametric – Capacity Computing Applied to Design Exploration

Because design for product integrity is a pressing engineering imperative today, engineering teams must progress from examining a single design point to exploring the numerous design points required for the widest range of material properties, manufacturing processes and real-world operating conditions. The greatest obstacle becomes the time required to run all the design points. Besides using parallel computing (as mentioned in previous sections), one way to reduce that time is to solve multiple design points simultaneously. To run those at the same time, however, multiple licenses also would be required, which makes running simultaneous design points cost-prohibitive for many users. ANSYS, therefore, recently introduced a new, more affordable licensing model that enables simultaneous execution of multiple design points while consuming just one set of application licenses. This value-based parametric pricing model is shown in Figure 5.

For parametric analysis, the ROI will not increase, in a linear fashion, with the number of design points as each were completely independent. For example, the value of the tenth design variation in a parameter study is likely higher than that of the hundredth design variation. Similarly the incremental value of this parametric HPC capacity decreases as the number of simultaneous design points increases (see Figure 5).

To perform parametrically linked simulations on multiple CPU cores simultaneously, ANSYS offers the following parametric HPC products:

- ANSYS HPC Parametric Pack: Parametric license for structural mechanics and fluid dynamics simulations featuring:
  - ***Automated parametric design studies:*** Ability to automatically and simultaneously execute design points while consuming just one set of application licenses.
  - ***Scalable cost model:*** Number of simultaneous design points enabled increases quickly with added packs (see Figure 6). In keeping with the law of diminishing returns noted above, the cost per design point drops as you increase the number of design points in a single study.

For example, enabling 64 instead of four simultaneous design points incurs a cost premium of roughly 3 times – and not 16 times that the design-point increase might imply.

- ANSYS Distributed Solve (DSO) with Optimetrics: Parametric license for electromagnetics simulations featuring:
  - ***Automated parametric design studies:*** Ability to automatically and simultaneously execute design variations while consuming just one set of application licenses.
  - ***Parametric design optimization & sensitivity:*** Ability to determine optimum design parameters as well as relative sensitivity of the output to the various design parameters.
  - ***Near-linear speedup:*** Offering near-linear scalability to a large number of cores.

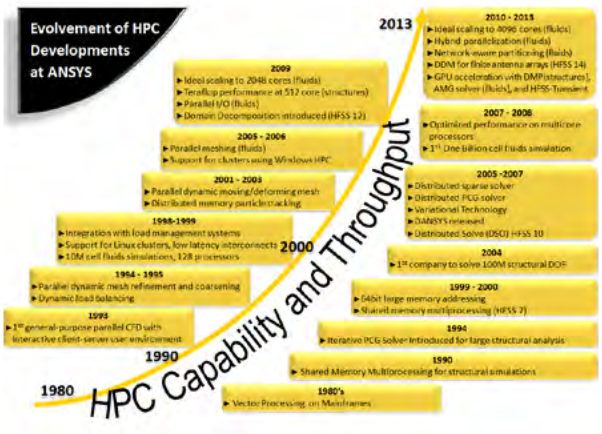

Figure 7: Overview of major HPC software advances made at ANSYS over the last three decades.

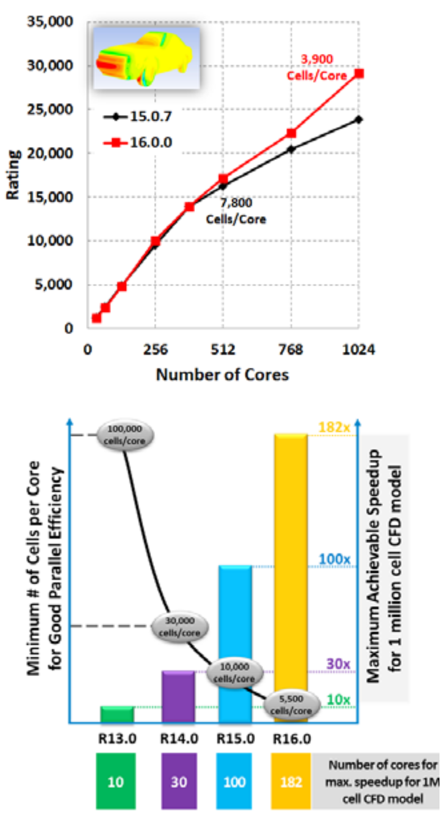

Figure 8: Recent improvements made in parallel scalability for a 4-million-cell model in ANSYS Fluent (top: a higher rating means faster performance). Bottom image shows the generalized effect of these release-after-release improvements made in terms of maximum achievable speedup for a 1-million-cell model.

Value of Sustained Investment in HPC

HPC has become an imperative part of all software development. As processor speeds have levelled off due to thermal constraints, hardware speed improvements are now delivered through increased numbers of computing cores. In turn, efficient execution on multiple cores is essential. Software development to build and maintain parallel processing efficiency is, therefore, a major on-going focus at ANSYS.

Figure 7 shows our history of delivering HPC technology, and reflects our sustained investment and achievement. Our HPC efforts – internally and with HPC hardware partners – lead to technology and performance leadership. More importantly, our HPC efforts ensure customer's return on overall HPC investment. This is not only for today but also for the future.

Thanks to HPC partnerships, the scalability and performance of ANSYS CFD have for example been improved release after release. As can be noticed from Figure 8, ANSYS can demonstrate good parallel efficiency at a lower number of cells per core, which implies that:

- All users will be able to decrease the time to solution (or better, insight) for their current models by leveraging more computing cores.
- All users will be less constrained by hardware limitations because bigger models can be sped up at existing computing capacity.
- Some– who have access to a large HPC system – could potentially benefit from groundbreaking scale-out performance at very high core counts for large models.

Because HPC is so dynamic, the computing landscape changes so quickly, it is critical that your engineering simulation software focuses on HPC so that you can take advantage of tomorrow's technologies and expand the scope of what you can accomplish. ANSYS is doing just that. Some examples can be found below.

- The first example is on-going optimization of our software. Here, we do much work with the processor vendors like Intel and NVIDIA to understand and optimize the mapping of the simulation workload to the compute resources (i.e. using new compiler options, setting processor affinities, and dynamically adjusting how the workload is balanced across multiple processors). ANSYS Mechanical customers can now interchangeably use CPU cores and multiple GPU's to accelerate most simulation workloads. In addition, GPU support for the (3D AMG) coupled pressure-based solver in ANSYS Fluent demonstrates our commitment to allowing customers to leverage new and evolving technology [1], such as GPU, for faster simulation. Finally, we are excited to be supporting Intel's Xeon Phi coprocessors in ANSYS Mechanical.

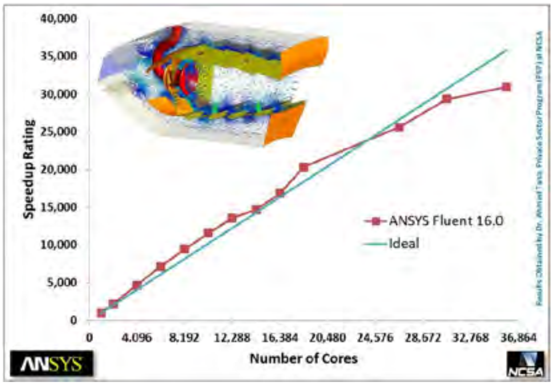

Figure 9: This image is from a benchmark [3] being run in collaboration with the National Center for Supercomputing Applications (NCSA). The combination of the Blue Waters supercomputer and the Cray Inc. engineering team, the NCSA Private Sector Program technical team, and the enhancements in Fluent 16.0 enables organizations to study the most complex and realistic simulation cases.

- The second example speaks to solution scalability and the future of HPC: how we will need to continue enhancing our software in order to take customer's use of simulation to the next level. ANSYS has partnered with Cray, Inc to deliver ground-breaking scale-out performance [2] at 14,000 cores for a large (105-million cell) industrial CFD model with ANSYS 15.0. While a milestone of scaling to 36,000 cores was more recently achieved with ANSYS 16.0 (see Figure 9), the goal of this ongoing partnership is to achieve parallel performance at 50,000 cores and more, and to continue to show great scaling as core counts per processor continue to increase.
- A third example is the extension of parallelism and HPC performance across all aspects of the simulation process – from meshing and setup to file handling, visualization, and automated optimization techniques. For example, we enable good scalability of parallel meshing when generating Tet/prism meshes. Although the performance is case dependent, 92% scalability has been observed on a 42 million cells mesh when using 8 cores.
- A fourth example is the usability of our HPC solution. We are driving toward a unified environment – across solver components – for defining, submitting, and monitoring customer's parallel workloads. In conjunction with our hardware partners, we are also exploring how our software can be more robust to hardware issues, and how we can find and resolve those issues dynamically, in order to avoid interruptions or failed jobs. Finally, another example of improving the usability of our HPC solution is through teaming up with HPC strategic partners (like Fujitsu, HP, and IBM) for simplifying cluster deployment. The delivered solution is a pre-configured, and ANSYS optimized HPC cluster, including support for remote 3D visualization and cluster/resource management.

Final Remarks

HPC is the present and future of engineering simulation.

Based on more than three decades of HPC deployment and associated software development, ANSYS recognizes the value and benefits HPC brings to customers. Therefore, ANSYS HPC software licensing is designed on pricing models that ensure the highest value for engineering simulation workloads while allowing ANSYS to continue our HPC software developments.

The value of ANSYS HPC licensing has recently been extended so that:

- More users can benefit from hardware accelerators (like NVIDIA GPUs and Intel Xeon Phi).
- Volume parallel processing is also rewarded at relatively low core counts for high-end workstations or entry-level cluster deployments.
- A virtually unlimited number of compute cores can be enabled for a single simulation.

More cores means more value and a competitive edge – and, yes, increased cost – but ANSYS ensures that the value of the HPC software exceeds the cost of investment that is required to stay ahead as HPC evolves into the future.

References

1. "Advances in Accelerator-based CFD Simulation", Wim Slagter, New Trends in CFD 2, September 2014, DANSIS Seminar.
2. "Cray and ANSYS Achieve Extreme Scaling Improvements on ANSYS Fluent Using CSCS's Cray® XC30™ System", Application Brief, October 2014.
3. "Supercomputing Milestones via Partnerships", Wim Slagter, ANSYS blog, November 2014.

ANSYS, Inc.
Southpointe
2600 ANSYS Drive
Canonsburg, PA 15317
USA
724.746.3304
ansysinfo@ansys.com