engineering.com

All You Need to Know About Hardware for Simulation

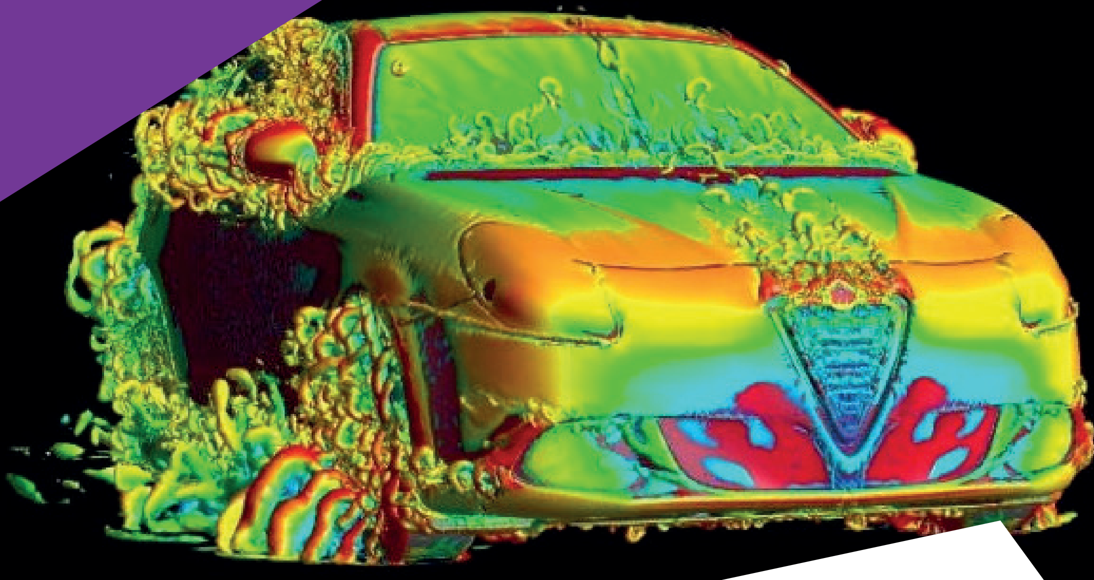

Ansys

intel

This research has been sponsored by Ansys and Intel.

TABLE OF CONTENTS

THE HIGHEST PERFORMING COMPUTING AND WHY ENGINEERS NEED IT

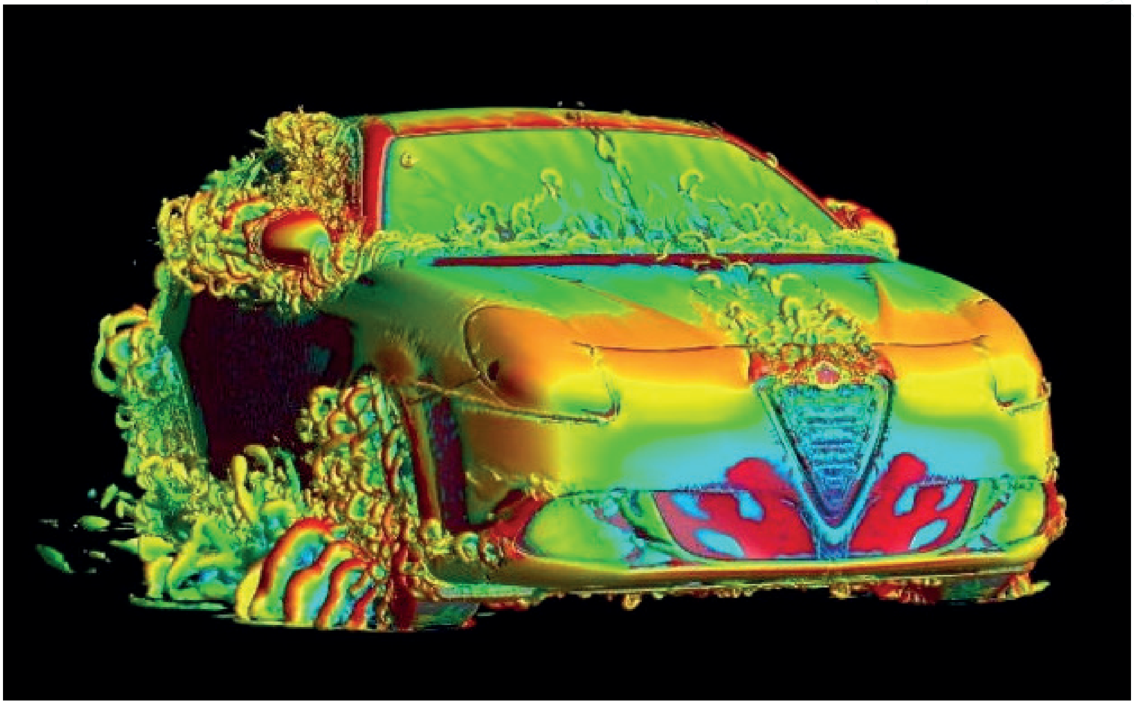

Predicting wind noise around the Alfa Romeo Giulietta with Fluent. (Image courtesy of FCA Italy.)

Compute power is available like never before. It may be local, under your desk or in your next workstation. An almost infinite amount of compute power is out of sight, in computers the size of pizza boxes; stacked in floor-to-ceiling racks; crammed with GPUs, CPUs and massive amounts of storage—all networked, all online, in row after row of racks in data centers spread as much as a quarter mile in each direction. These data centers are all over the world—over 1,800 are located in the U.S. alone. They are popping up in cornfields and deserts, wherever there is space—and power.

This vast scale and number of data centers make up cloud computing, and in the cloud is where your split-second search and video streaming—as well as your project collaboration, file storage and, for the purposes of this article, where your engineering simulation—already lives, or will do so soon.

In the engineering and design world, compute power is never enough. The ideas that there is unlimited computing power on high performance computers (HPC) on cloud networks is tantalizing.

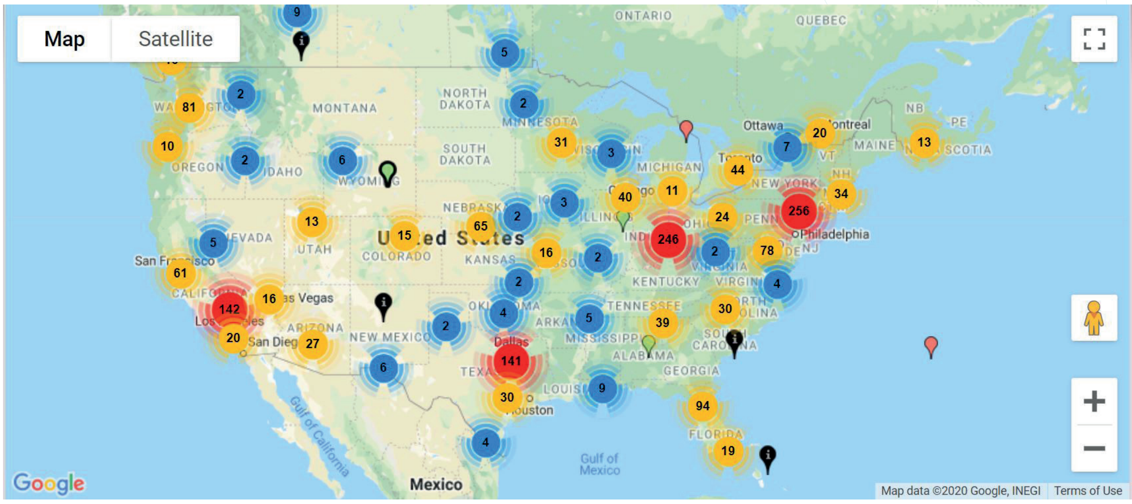

Data centers are popping up everywhere. Compute power may never be enough for engineers, but with over 1,800 data centers located in the continental U.S. alone, each crammed with compute power and storage, we are living in an age where compute power is cheaper and more plentiful than ever. The San Francisco Bay Area alone has 61. (Image courtesy of Datacentermap.com.)

ISN'T A WORKSTATION ENOUGH?

For engineers in startups, consulting practices and in small engineering firms, simulation is done on local computing resources, in most cases a personal workstation. These workstations can be limiting; but for many who are comfortable with their workstations, anything else—like HPC—raises questions and concerns. How can we leave our workstation? What would we look for in an HPC configuration? How should HPC be deployed? How much will HPC cost, and can we justify the expense?

Understandably, engineers will need some reassurance before they let go of their trusted workstations. The engineer's fascination for top-of-the-line gear may do the trick. A supercomputer is like a supercar; a Ferrari compared to the practical Toyota. An engineer may not be able to get a supercomputer—but they can use HPC can get supercomputer-type performance whenever they want.

With this report, we will explore all the hardware choices available for simulation as well as try to demystify HPC. You will learn how HPC can useful, accessible and affordable for everyone—not just big companies. Read on and you will discover:

- What is HPC?
- Why engineers need extra-strength computing.
- Myths about HPC.
- Saving time with HPC.
- How HPC can pay off.
- What to look for in a workstation.
- HPC hardware and configurations.
- Buying or "renting" HPC.

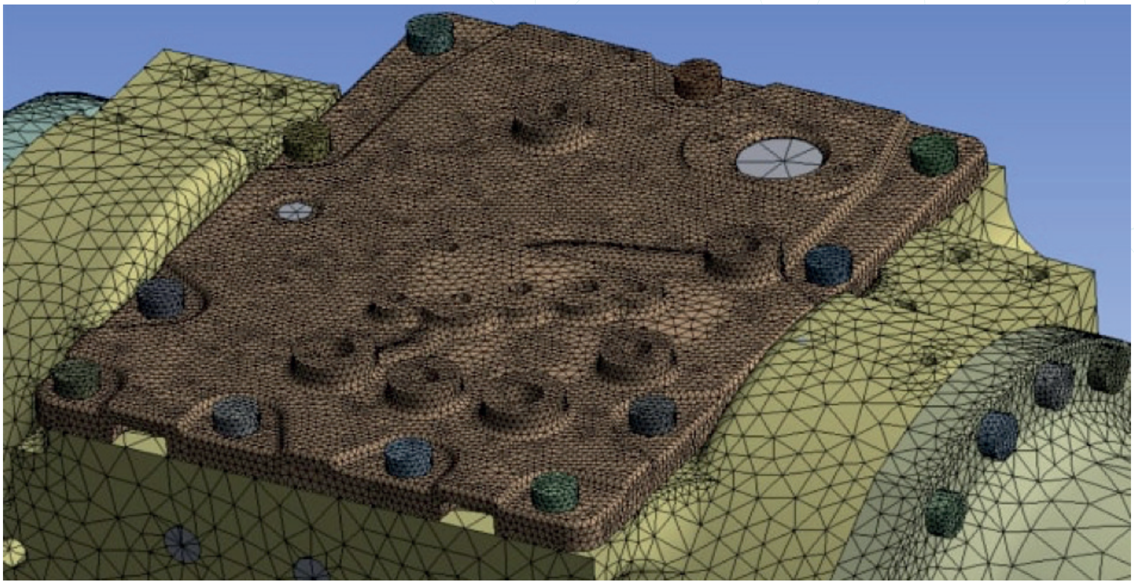

Your CAD/ CAE software is not doing you any favors by providing a million-element model that will take your workstation forever. Engineers spend time "defeaturing" a finite element model to reduce solution time. But HPC can solve large problems fast, which means engineers don't have to defeature. (Image courtesy of Ansys.)

INFINITE ELEMENT ANALYSIS

The most common engineering simulations, finite element analysis (FEA) and computational fluid dynamics (CFD), solve systems of equations based on the nodes of a mesh. The meshes are approximations of the original NURBS-based (non-uniform rational basis spline) CAD geometry. The finer the mesh, the more fidelity it has to the geometry and therefore will yield the most accurate results. But as element size shrinks to infinitesimal, the number of elements and degrees of freedom approach infinity. Even a simple part can be meshed so fine that it will take hours on a supercomputer.

So, if your simulation is not limited by hardware, you are not trying hard enough.

Early on, the engineer doing simulation learns to compromise. Let's not model this detail, or that one. But with every detail not modeled, the greater the deviation from the designed part and the more erosion of the model's fidelity, the greater the approximation. But what can you do? You need to take the load off the hardware.

"Engineers are constrained by turnaround time limitations," said Wim Slagter, director of HPC and cloud alliances at Ansys. "So, they are spending time and effort on changing their model in order to be able to run it on their existing hardware or to get acceptable runtimes. They must make their models smaller. They are trading off accuracy of results by reducing the number of elements, the number of features—or using a less advanced turbulence model in their CFD application just to get acceptable runtimes."

Also, consider the opportunities lost because of the limits of local computing. The transient and turbulent problems that won't get done, the fatigue that is not explored, nonlinearities not looked at and thousands of generative designs never generated – all because the computing resources you have on hand can't handle it. We can't all have supercomputers.

MOORE'S LAW IS NOT ENOUGH

But are supercomputers, at the cost of tens of millions of dollars, the way to go for today's engineering simulations? An alternative to supercomputers is tying together ordinary computer hardware in massive networks to create what is known as high performance computing (HPC). The economies of scale of HPC have made solution times faster and cheaper than they have ever been—and circumvented Moore's law.

Gordon Moore, once CEO of Intel, famously predicted in 1965 that there would be an annual doubling in the number of components per integrated circuit. This has come to be known as Moore's Law.

Moore's law has made for smartphones that are more powerful than the computer that put Apollo astronauts on moon. And mobile workstations now run circles around your dad's minicomputer. But the wealth of computing on the cloud is faster by a long shot.

Gordon Moore may not have seen HPC networks coming, which take advantage of parallel processing, with each processor handling a calculation. If the software is able to feed each processor a calculation, all of the calculations can happen at once, in parallel. Even with ordinary commodity processors, running in parallel can be faster than sequential processing, where one calculation occurs after the other—even if they are done on a supercomputer.

WHAT IS HPC?

If you are queuing up a simulation today on your workstation, it will be ready tomorrow; the same solution on HPC on the cloud will be done by lunch. The cost for simulation is fractions of a cent per second. You might wonder why we should buy any local workstations at all. Why not just plug into the massive grid of unlimited speed and power offered by this new world of computing and data centers?

This is what computer hardware looks like in the big leagues of simulation. Pictured: The High Performance Computing Lab at George Washington University. (Image courtesy of Rackspace.)

HPC systems came into play in the early days of Bitcoin mining. Today, HPCs are used in brick-and-mortar establishments for some serious problem-solving. In engineering, HPCs are used for complex simulations, including transient problems, acoustics, CFD and combustion in engines.

In multi-physics simulations, the simultaneous effect of multiple states of matter is studied, including gas, liquid and solid, as well as phase changes. An example is analyzing how the effects of thermal expansion affect a moving structure, as well as the turbulence caused by the moving structure over several time steps. Here, paired partial differentiation and other complex mathematical modeling must be calculated to solve the problem. HPCs can speed up this process of simulation.

In design and manufacturing, HPC has been used to solve structural and thermal design problems, and to optimize production lifecycles.

Special networking software appropriates the resources it needs from among the networked computers, allowing faster solutions of large and complex problems that a personal workstation would not be able to solve in a reasonable time, or at all.

HPC can be inside your facility—“on premise”—or outside your company and, as is very common these days, in the cloud.

In a recent survey sponsored by Ansys, out of over 600 engineers and their managers, about a hundred of them (17 percent) were using the cloud for engineering simulation. But an additional 20 percent were planning to do so over the next 12 months.

Is an HPC configuration a supercomputer? In strength and output, perhaps; but architecturally, there is a fundamental difference. An HPC configuration is an aggregation of many computer cores in separate computers which are connected by an operating system, whereas a supercomputer is a monolithic, standalone computer, all in one spot.

The early Cray computers decades ago were standalone supercomputers. These computers were singularly equipped with the customized hardware needed to handle complex simulation and calculations. Today, Cray computers have embraced the HPC concept and offer clusters of supercomputers, the Cray CS (for cluster supercomputers), raising the high end of HPC configurations.

Other important distinctions to consider when comparing and contrasting supercomputers and HPCs is cost and customization. Most supercomputers run only certain customized software applications, and not every engineering simulation software familiar to engineers may be compiled for a supercomputer.

Then there is the cost, which can be considerable—up to tens of millions of dollars. The cost comes not only from acquisition of a supercomputer, but the cost of running and maintaining it. Cost alone puts supercomputers out of the reach of most engineers. HPCs on the other hand, are interconnected affordable computer systems that can run legacy software, and are scalable by adding even thousands of low-cost computer "nodes." It is this sheer number of nodes that gives HPC the computing power of a supercomputer, without the cost.

HOW HPC WORKS

Every computer in an HPC system is known as a node. Each node is generally equipped with multiple processors called compute cores that handle the computation aspect of problem-solving. The processors, graphical processing units and memory of each node are then interconnected by a network to make a high-performance computing system.

You can compare HPC systems to the rendering farms used for movie special effects and realistic architectural fly-throughs. Multiple compute nodes work together to deliver solutions for large, complex problems. When multiple processes line up for an HPC system, a scheduler comes into play. The scheduler allocates compute and storage resources to each process according to its requirements.

Approximately 64 percent of HPC platforms now integrate cloud computing and the market is expected to grow annually by 12 percent within the next 5 years.

CORES USED AS NEEDED

HPC is good for simulation because a large number of cores can be used for calculation. A core, or CPU core, is like the brain of a CPU. CPUs can have multiple cores. The Lenovo ThinkPad P1 mobile workstation CPU has 8 cores. An HPC configuration can have hundreds of cores.

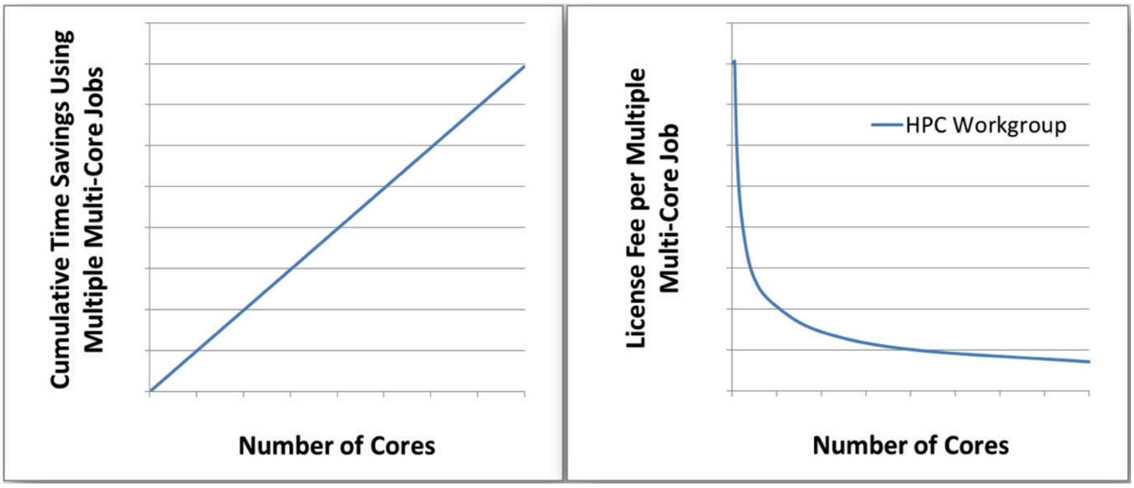

Running multiple multi-core jobs. The cumulative time savings increases linearly with the number of cores used. (Picture courtesy of Ansys.)

Above, we see a scenario where an Ansys customer is running multiple multi-core jobs. Here, the cumulative time savings increases linearly with the number of cores used as multiple jobs are executed simultaneously using HPC.

“This is what I call ‘HPC capacity computing,’” says Slagter. “Faster turnaround can be achieved, which linearly increases simulation throughput and productivity. For this we have HPC Workgroup licenses through which we are rewarding the volume buyer with scaled pricing. Our HPC Workgroup licenses provide volume access to parallel processing for multiple jobs and/or multiple users.”

“Even software with an ideal, linear scalability like our CFD solution, the time to solution asymptotes,” says Slagter. “So, there is clearly a diminishing value of added cores after some point. You may start saving minutes instead of hours, or hours instead of days. Our value-based pricing model accommodates this idea of decreasing incremental value of parallel as the number of cores increases. Our customers pay less as they add more HPC capability.”

Ansys users who need to take advantage of HPC will receive an estimate of the cost of cloud-deployed HPC up front from inside Ansys, according to Slagter. “We are quite transparent that way.” Users will also get an estimate of the turnaround time before they submit the job.

MISCONCEPTIONS ABOUT HPC

HPCs have predominantly been used by big businesses, which has led to various misconceptions about its use or applicability. Perhaps the most common misconception is that HPC can only be used for super-complex simulations, particularly CFD. This may be due to the scale of CFD simulations and the underlying equations. In reality, HPCs are suitable for almost any type of simulation, including structural and electromagnetic simulations.

Cost-conscious management will always offer a sobering thought when HPC is brought up: that HPC is too expensive for a small or medium sized enterprise, (SME). Most project managers and company executives still believe HPCs is high-priced and offers too low of a return on investment (ROI). The boss who balks at buying a $5,000 workstation every five years will not be keen to buy an HPC configuration. Who knows how much that would cost?

But a strong case can be made for the economics of HPC with the advent of cloud based HPCs, covered in the next section of this report. For now, suffice it to say that the cost of using HPC has come way down. Simulation vendors use cloud services from Amazon, Google and Microsoft to offer scalable pricing models at a low cost for complex engineering tasks. Expensive computers don’t have to be purchased—so there is no budget hit, and no panicking accountants. You simply rent the computers and pay for what you use. You can do a serious simulation for the price of a large latte.

The last myth is centered around ease of use. Many believe a dedicated IT supervisor or manager is needed to handle HPC-related activities. However, the maintenance and upkeep of an HPC equipment can be someone else’s headache; HPC systems can be managed offsite, and you never see them or have to deal with them. Companies can also contract the HPC reseller to take care of the HPC sold to them.

THE BUSINESS CASE FOR HPC AND THE CLOUD

You might value the perfect simulation model, striving for perfection—an engineering ideal if ever there was one.

"Perfection is the enemy of the good." Your manager misquotes Voltaire as he studies his engineers' timesheets. Was all the time spent simplifying the model worth it? The manager knows that time is money, but why do the engineers not care? They spend hours or days defeaturing, or they obsess over tiny cells in a fluid model where the flow is not critical, and they manually adjust meshes.

They do that in their search for perfection, but also to make sure the model doesn't blow up their workstations, that it fits in storage and memory and that it solves in time. The clock is ticking, and time is money.

IS SPEED WORTH IT?

While computing power on workstations feels limited, the amount of computing power available on the cloud is downright staggering. So why not use the cloud and its unlimited capacity to blaze through an FEA or CFD model—even one with millions of elements or cells—rather than have engineers grind away dumbing down the geometry.

Using the CAD geometry as is, rather than simplifying it, cuts a time-consuming step out of the process. So what if every detail is meshed? It costs pennies per hour – much less than the cost of engineers. Let your engineers do more simulation rather than whittle away at a CAD model. Let them go ahead and solve transient and turbulent problems, explore fatigue, go nonlinear, explore a dozen alternate designs—all worthwhile uses of high-priced engineering talent.

RENT OR BUY HPC

While big firms in the aerospace and automotive industries can justify buying an HPC configuration, build facilities for them and maintain it all with an IT staff, small and medium-sized companies will do well to rent HPC, rather than buy it. Renting may still make sense for big companies when they run out of conventional computing resources, as Rolls-Royce did and which will be discussed later in this report.

Amazon Web Services (AWS), Google Cloud Platform, Microsoft Azure, Nimbix, Rescale and Penguin Computing On Demand are some of the leading vendors for using HPC on the cloud on an as-needed basis.

Amazon, the HPC leader according to Cloud Lightning, will create a virtual cluster of HPC hardware for you to run your next job. AWS lists CFD as one of its main HPC applications. You can spend $0.75 an hour trying your application in the AWS sandbox.

It may be uncertainty about the final cost of analysis that keeps users from taking the plunge into HPC. It would be nice to know how much an analysis will cost. First-time users may be tentative, and skeptical of unfamiliar HPC providers that could run up the bill.

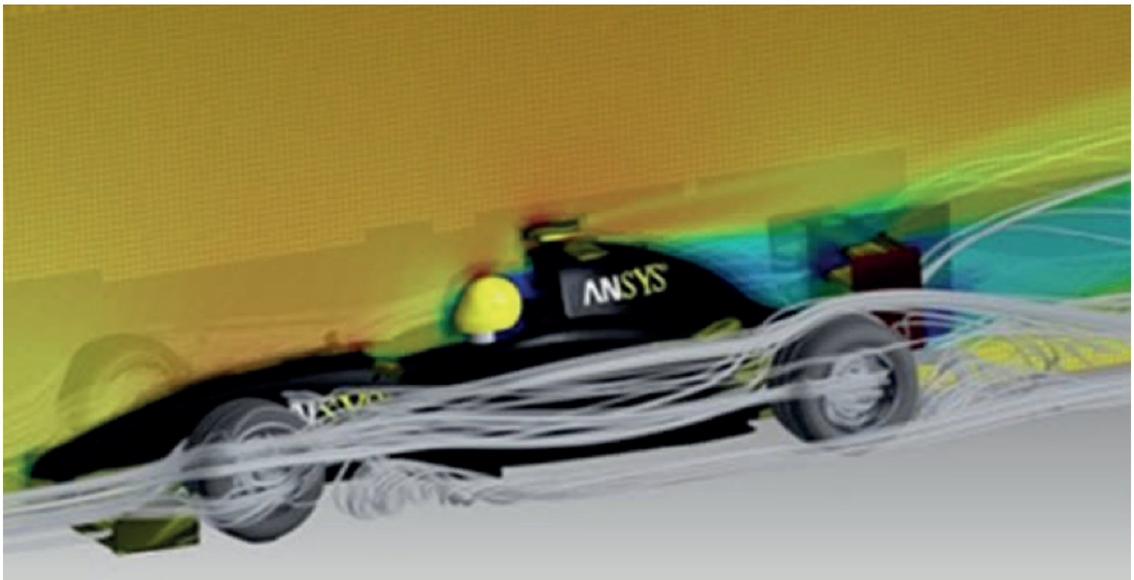

F1 race car simulation on HPC. (Picture courtesy of Ansys.)

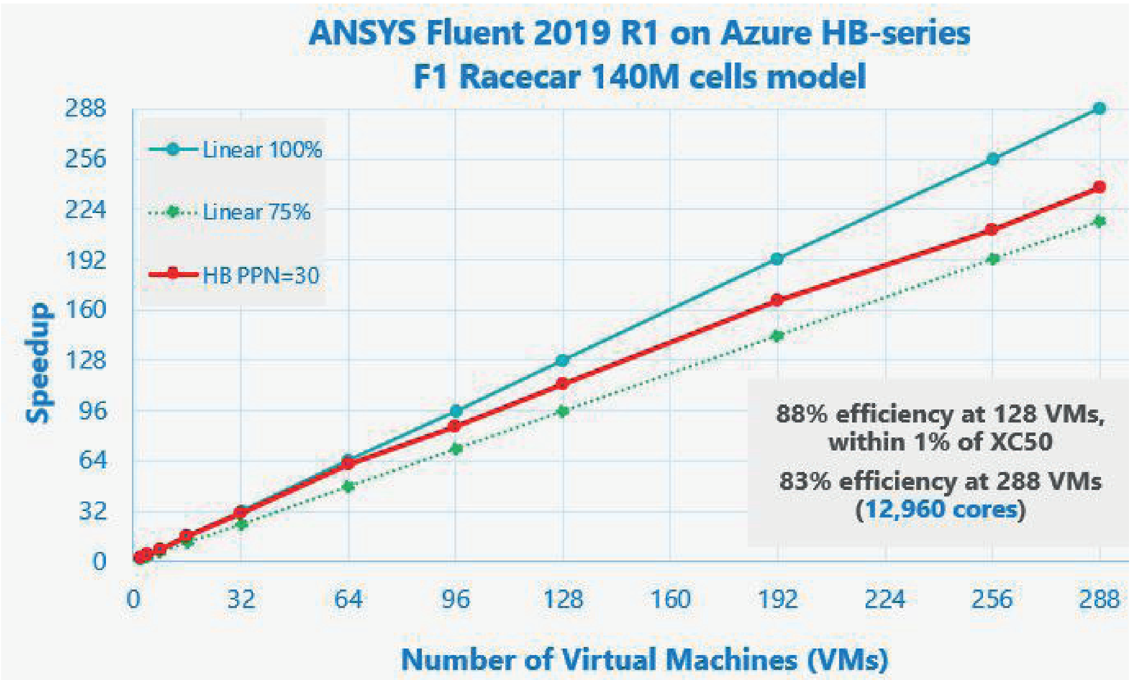

HPC service providers do provide online calculators to predict solution time. However, the final cost in dollars is not usually public information. We estimate that, if wall time is equal to CPU time, the 140 million-cell CFD simulation above with a 128-core setup from Microsoft Azure would take 1,850 seconds and cost under $4. The cost on AWS would be similar.

GOING TO THE CLOUD FROM YOUR APPLICATION

Ansys offers two ways to use HPC on the cloud if you don't want to buy your own HPC hardware. From within several core Ansys applications, you can pick Ansys' own HPC service, the Ansys Cloud, on which to run the solution. Or, you can choose one of more than 10 Ansys cloud hosting partners.

Switching to a cloud solver from within the application seems to be the easier approach.

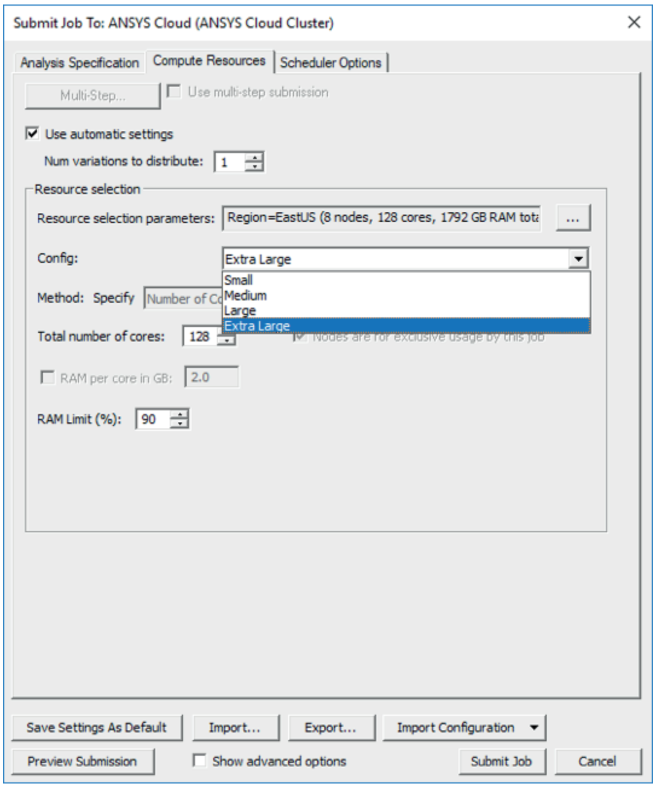

When the workstation is not enough. Dialog box for doing the calculations on the cloud. (Image courtesy of Ansys.)

Let's say somebody is using Ansys Electronics Desktop and they realize the problem is too large or will take too long.

"Once they reach the limits of their desktop computer, they can basically select a menu and start selecting from a few standard configurations on the Azure cloud," explained Slagter.

The Ansys user can select small, medium, large and extra-large cloud-based server configurations that correspond to 8, 16, 32 and 128 cores, respectively.

"Users don't have to go to Azure directly, or set up an agreement with Microsoft," said Slagter. "It's all handled by Ansys. They pay through a single license key for software and also for cloud hardware cycles."

"It is extremely easy on Ansys Cloud because it is our own solution," continued Slagter. "We have developed that user interface, whether it is with Ansys Mechanical, Ansys Fluent or Ansys Electronics Desktop. If users reach the capacity limits with their on-premise computers with a computationally demanding job, they can easily switch to the cloud and run on it. That is really easy because we developed a seamless interface to the Microsoft Azure cloud.

"They can post-process it on the cloud, too, so there's no need to immediately download the results. They can do some lightweight reviewing of the simulation results in the cloud and if they want to do in-depth post-processing steps, they can transfer the results back to their desktop machine for further investigation."

"If customers want to use HPC on a partner-managed cloud solutions provider, there may be a few more steps required," explained Slagter.

PARTNERS PROVIDE

If Ansys users want to work directly with companies that provide HPC services, Ansys provides a list of HPC vendors that offer streamlined HPC services to companies, small or large.

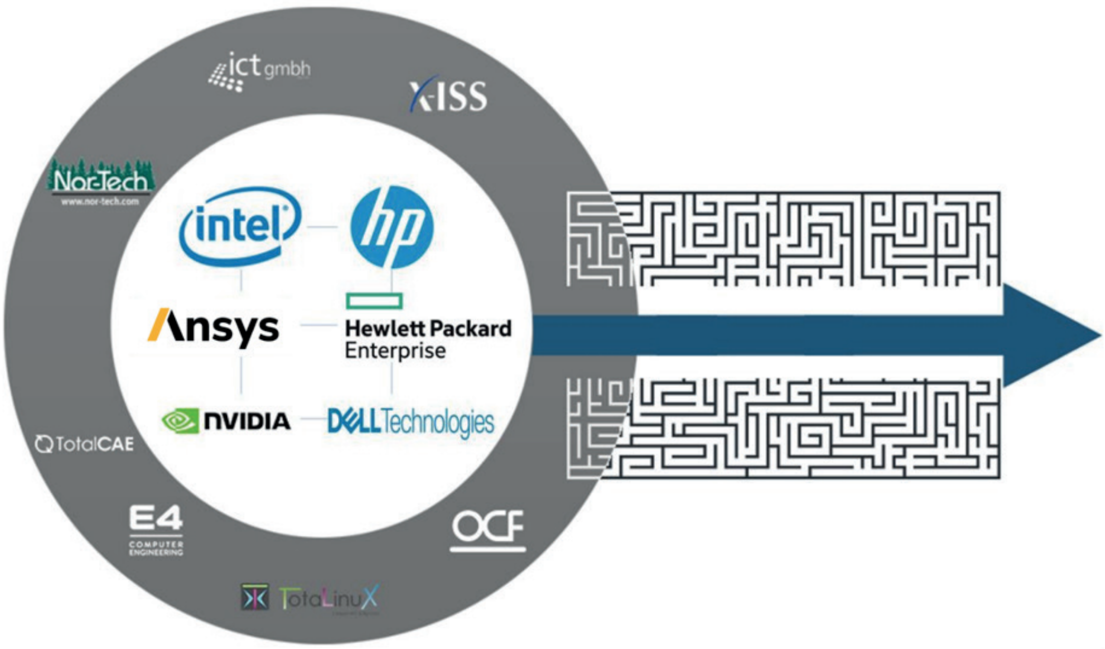

Ansys works with many HPC strategic partners to make specification and deployment easier for their customers. These vendors also offer post-sales services to ensure integration and deployment tasks are properly executed. (Image courtesy of Ansys.)

"They have proven their solutions, they have developed their own simulation environment for Ansys which is compliant with our platform requirements, and of course they also make sure that it's as easy as possible," said Slagter. "Of course, it can't be as easy as it is with the Ansys Cloud, but for customers who need to have HPC on premise—a significant portion of our userbase—we have an option for them."

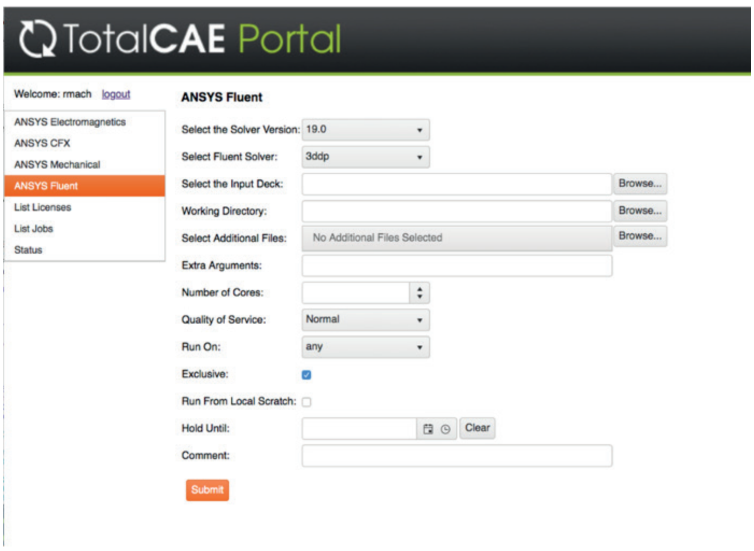

The TotalCAE portal will let you select a model file and submit it to their HPC hardware. (Image courtesy of TotalCAE.)

While access to the HPC partners is not from within the Ansys software, some partners do tailor their applications to make it easy for Ansys users. For example, TotalCAE has a CAE portal that lets you select your model file with the push of a button, it submits the job to their HPC hardware. TotalCAE manages the entire HPC appliance and Ansys portfolio for the customer, basically removing IT as a barrier to adopting HPC. The company makes no mention of pricing, however.

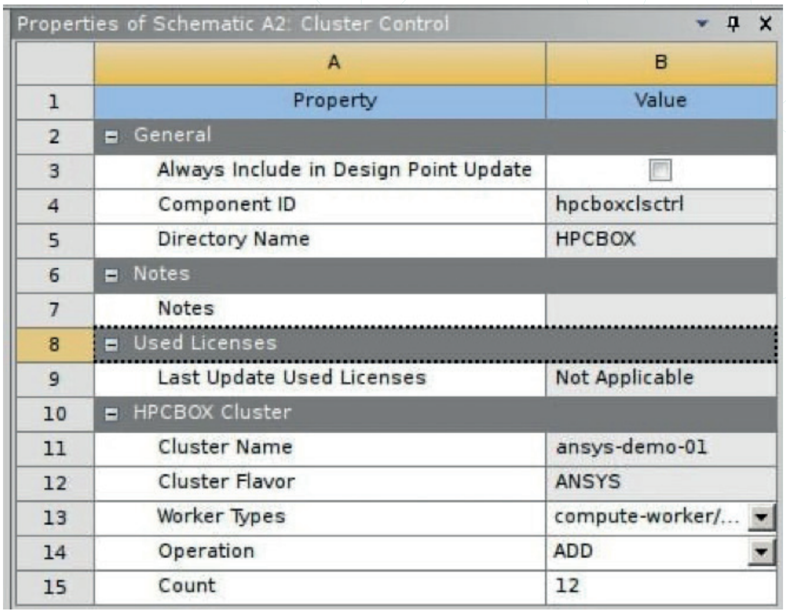
Properties of Schematic A2: Cluster Control

	A	B
1	Property	Value
2	General	
3	Always Include in Design Point Update	
4	Component ID	hpcboxclsctrl
5	Directory Name	HPCBOX
6	Notes	
7	Notes	
8	Used Licenses	
9	Last Update Used Licenses	Not Applicable
10	HPCBOX Cluster	
11	Cluster Name	ansys-demo-01
12	Cluster Flavor	ANSYS
13	Worker Types	compute-worker/...
14	Operation	ADD
15	Count	12

HPCBOX by Drizti HPC menu. (Image courtesy of Drizti.)

One very Ansys-centered HPC partner is Toronto-based Drizti with its HPCBOX. Like TotalCAE, Drizti also manages the software licenses, and users are not required to create even one line of code. HPCBOX, while not push-button easy, uses a "Connector" interface that appears to provide more control. Drizti also makes no mention of pricing.

Rescale, an eight-year-old San Francisco startup, claims to have the largest HPC infrastructure network in the world, with access to eight million servers. This is remarkable growth and may be a testament to the need for a total service, one that lets you rent engineering solutions as well as the HPC hardware the solutions run on. Rescale does let you use your software license, if you have one. You can drag in an input model or select one. A vast menu of simulation software includes Ansys, Siemens, Dassault Systèmes, MSC, OpenFOAM and many more.

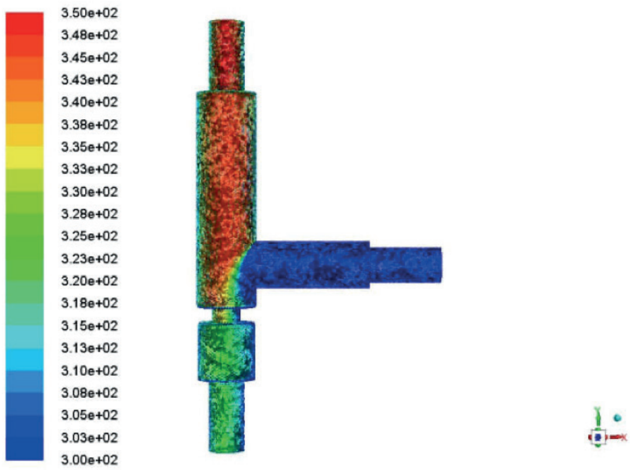

Ansys Fluent T-junction mixing example done with Rescale, a leading HPC cloud service provider for engineering simulation, shows a hot water and cold-water inlet, with mixing and calculation of temperature at the outlet. (Image courtesy of Rescale.)

The list of engineering software supported by Rescale is so long you will wonder what engineering software is *not* supported. Rescale interactively lets the user pick the processors, number of cores, etc., and displays the price of the simulation. For example, a quarter hour with 144 cores would cost less than $5.

HPC: IT'S ABOUT TIME

To perform a simulation, you start with a 3D CAD model and make a mesh; however, only the greenest—or laziest—engineer would take the entire CAD model and make a mesh from it. Many details on a CAD model don't matter to the analysis, and the mesh will be too big and solving it will take too much time. Useless details must be removed in a process called *defeaturing*.

Defeaturing costs a lot valuable engineering time.

But even after defeaturing, there are millions—if not billions—of finite elements (for a structural simulation) or cells (if it is a flow simulation). The diligent engineer, by tradition, will now go to work. They must sort out how the nodes are to be connected, how contact surfaces are to be established and how meshes can be refined. The engineer must always be conscious of the number of degrees of freedom. Each node represents up to 6 degrees of freedom (DoF), and each DoF requires a calculation. Collectively, these calculations can overpower a supercomputer. You may not get any solution, much less multiple solutions, in the time you have available.

DEFEATURING IS DEVOLVING

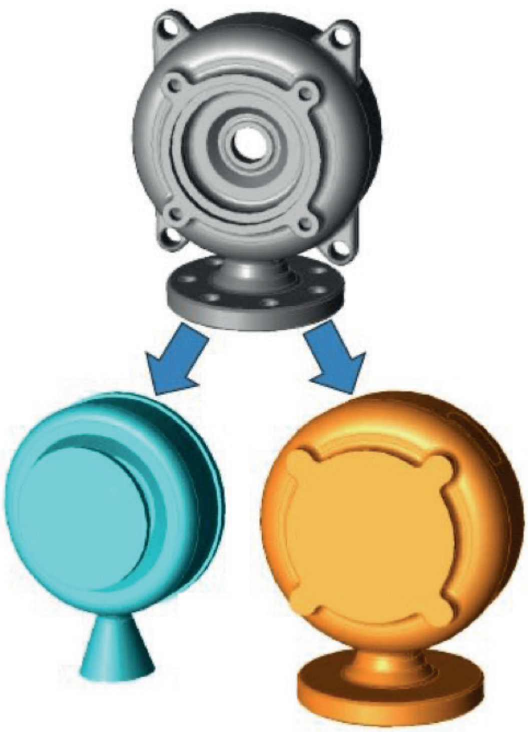

A CAD model is "defeatured," a manual or semiautomated process that seeks to reduce the size of the mesh to be used for simulation. Do it well and details that do not affect simulation will be safely removed. However, each defeature departs from the fidelity of the mesh. Do it badly and the simulation will avoid critical areas. (Image courtesy of TransMagic.)

The urge to see how fine a mesh you can get away with is constant. How fine can you go before your results don't come back in time?

To get to a manageable problem size, engineers pull out all sorts of tricks. A fine mesh in critical areas, a coarse mesh in others. We agonize over how small our elements can be, how fine the mesh can be and we stick to lower order elements. We model materials as homogenous, ignoring non-homogenous properties at the microscopic scale. We assume a steady state because time steps involve multiple solutions. Who has time for that? A fillet will result in a fine mesh, so it is reversed into a sharp corner. In the best case, you have reduced the number of elements, cells, degrees of freedom—all without affecting the outcome. The stresses or flows do not change in all the areas you care about. Sure, it took a little time, but it was worth it, right?

So we go around a part or volume removing chamfers, fillets, protrusions, holes and other details. We can set the minimum size to reduce the number of elements. Although CAE tools often claim they can fully defeature a part—and it may be fun to see all the tiny holes and fasteners disappear from your model—veteran analysts have learned not to avail themselves of wholesale push-button automation out of caution. Instead they resort to defeaturing more or less manually, which is a time-consuming and laborious effort.

We remove detail, mesh, count the elements, then repeat. Each time you mesh, the number of degrees of freedom gets lower. Another wave of defeaturing. All that so the model will solve in time – or solve at all.

Engineers see defeaturing a model as a necessary evil, and it's become a learned skill. Careless defeaturing can reduce the accuracy of results. In one study of a part where one feature after another was turned off, the accuracy of results (with a fully featured part representing 100 percent accuracy) dropped to as low 32 percent.

In a recent survey done by Ansys and Intel with 1,800 engineers responding, 75 percent felt forced to limit the size and amount of detail in simulation models just to complete the solution in the required time. However, almost all of respondents (87 percent) are not happy with the results.

How often do you limit the size and amount of detail in simulation models due to turnaround time limitations?

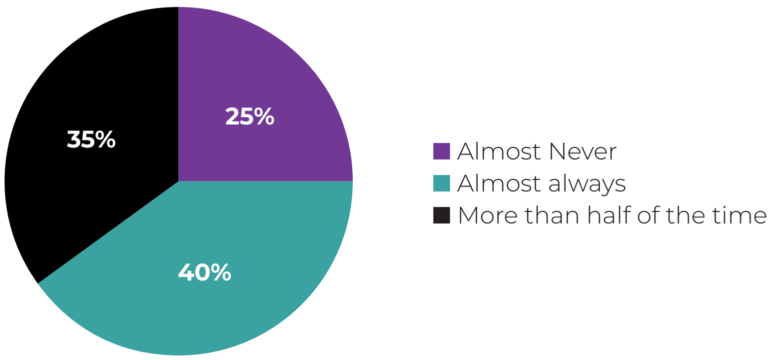

How often does "defeaturing" produce less useful results?

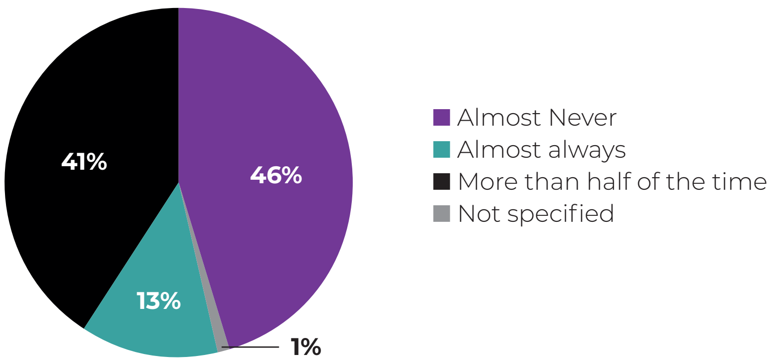

"Of course, engineers are computer-bound. They are also constrained by turnaround time limitations," said Slagter. "That became very clear through our survey. About 40 percent of the respondents limit the size or amount of detail for nearly every simulation model. 35 percent of the respondents limit it more than half of the time."

"That means that customers and engineers are constrained by turnaround time limitations because they are spending time and effort on changing their model in order to be able to run it on their existing hardware, or to be able to squeeze it on their existing hardware, or to be able to get acceptable runtimes," Slagter added. "So, they are actually making the model smaller, or maybe having a trade-off on accuracy because they are reducing the number of elements or the number of features—or maybe they take a less advanced turbulence model in their CFD application, in order to get acceptable runtimes."

DEATH BY DEFEATURING

As we work our way around the CAD model, removing details so the mesh will fit in our computers and in our schedules, we cross our fingers hoping the detail we remove is not critical—like that bolted joint.

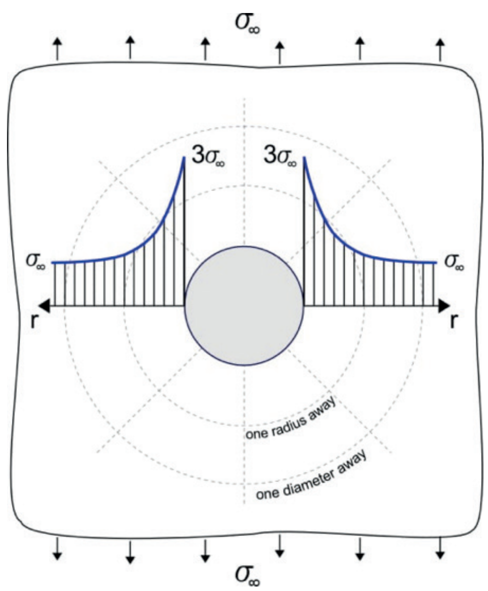

The stress concentration factor around a hole is 3X. Let's hope that hole was not "defeatured" in an attempt to reduce problem size. (Image courtesy of Fracture Mechanics.com.)

In the worst-case scenario, removing detail by defeaturing can be dangerous. You might have glossed over a critical area and ignored a potential cause of failure.

Let's say you follow a routine procedure as prescribed by FEA wisdom and you remove all features less than a certain size from the CAD model—so all holes of that size or smaller disappear. What if one of those holes is in an area of high stress? The stress concentration factor at the 3 o'clock and 9 o'clock positions around a sample in pure tension is 3. What if that would have produced stresses in excess of a failure criterion?

COST OF DEFEATURING

Let's say a simulation takes a hundred hours, which includes getting the model from the CAD program to the final report. In rough orders of magnitude, let's say the first step of reducing the unnecessary complexity of the model by defeaturing takes a tenth of that number, or 10 hours. By all accounts, this may be a conservative estimate, and with an engineer's time worth a $100 an hour (again, in rough orders of magnitude, one significant digit precision, but bear with me), that comes to $1,000 per analysis. A full-time analyst can be expected to perform 10 of these per year, which makes for a cost of $10,000 in labor for defeaturing per analyst.

LIMITED BY OLD HARDWARE

While computer service providers can be expected to have the latest, fastest hardware, the same cannot be expected of an engineer who is forced to use old hardware due to capital expenditure restraints. There is quite a bit of angst with mesh size, in terms of adjusting models to have them fit within the capacity of your aging workstation.

“We found out that a large portion of engineers run on relatively old hardware,” noted Slagter, citing a survey Ansys did with Intel. “We know they could benefit from recent hardware, the latest processor technology. Our software takes advantage of the latest Intel processors, for example. People can improve their performance using those recent Intel processors tremendously. And by the same token, simulation models are getting bigger in terms of size and complexity. Engineers have to run an increasing number of design variants to ensure robustness, for example. And we see also an increasing number of engineers who are computer-bound or constrained by the lack of computer capacity.”

To find out how long it takes to recoup investing in a new workstation, Ansys offers an ROI calculator.

TIME TO MARKET GOES DOWN, COMPLEXITY GOES UP

The world is getting more complex, the drumbeat of technology gets quicker and products appear faster.

In a study of product development over 15 years, product cycles were found to have been shortened by about 25 percent across all industries. In fast-moving consumer goods, the design cycle is half of what it was in 1997, according to a 2012 study of European companies by Roland Berger consultants. During the same period, the number of products in the same industries increased by three times.

Simulation has been getting more sophisticated in order to keep up. The previous generation of analysis may have had one set of engineers doing structural analysis, another set doing aerodynamics and yet another studying acoustic effects. Now those processes may all be combined in a multiphysics environment. This puts an increasing strain on simulation teams and their hardware.

“The complexity lies in the details and in phenomena that exist on a small scale—which can be transient in nature, like acoustics models or combustion phenomena or multi-physics applications,” confirmed Slagter.

“The other reason is that markets expect more innovative products that come to market faster,” he added. “To make things more difficult, we have requirements that conflict with each other. In order to differentiate products from the competition, engineers will need to explore more design iterations, do more trade-off studies, make course corrections quickly—all leading to new designs that can win against the competition.”

With an increase in the amount and sophistication of analyses as the result of market pressures, engineers bump up against the limits of their computers.

“Engineers may also be looking at various manufacturing processes and multiple operating conditions,” said Slagter.

Ideally, to save time, all the design variations could be run in parallel. However, that would require a workstation and a CAE software license for each one—resources not available for most organizations. HPC simulation can get around this problem as all the design variations can be sent at once and the solutions run in parallel, on multiple cores, or what Ansys refers to as an “embarrassing parallel computing.”

They offer a “more affordable licensing model,” according to the company’s white paper, that solves multiple design variations with only one set of application licenses. In effect, this removes the insurmountable barrier of one license/one workstation per solution and enables the vast time savings of running all design variations simultaneously rather than sequentially. The little additional cost per each added design variation only increases the ROI of HPC.

SIMULATE VS. TEST

Simulation can reduce the time required if it replaces physical testing. For example, consider the amount of testing that is required before autonomous vehicles are certified. Billions of miles must be driven in test vehicles. Even with a fleet of a thousand cars, there are still a ridiculous number of combinations of environments and situations to be cycled through. By the time all the testing is done, no one reading this will be alive. With every year, another forty thousand people die in human-driven vehicles in the U.S. alone. That should be enough of an argument to shorten the testing cycle using simulation, the only safe way to replace testing.

A proper simulation can be done in a fraction of the time it takes to perform real road testing. The billions of miles of testing can be driven at breakneck speed in simulation, in parallel on parallel processors. Mistakes made can be corrected without risk to humans. With technology under the public microscope, and with zero public tolerance for even a single catastrophe, fast simulation may be the only answer.

HOW MUCH FASTER, EXACTLY?

How much faster can a simulation be compared to simulation done on the aging beast of a workstation located under your desk?

The time savings afforded by multiple cores can be dramatic at first. If a CFD simulation takes 16 days on a single core, adding 32 cores might reduce the simulation time to four hours -- slashing more than two weeks from the schedule. Doubling the number of cores from 32 to 64 might cut the time in half again – from four hours to two hours. However, with further increasing cores, the effect is less dramatic, until continued increases provide only incremental gains measured in minutes or seconds.

"The incremental value of HPC decreases as the number of cores increases," confirms Slagter. "We have a value-based pricing model to accommodate that."

You might expect the cost of employing an increasing number of cores to be proportional to the increasing number of cores. That is a misconception of HPC, according to Slagter. Running 2,000 cores instead of 20 incurs a cost premium of only 1.5 times and not the 100 times that the core increase might imply,

Flow simulation is arguably the most computer-bound of all engineering simulations. CFD models are composed of millions, if not hundreds of millions of cells, leading to overnight runs on local workstations – if they will run at all. This might easily be the orderly world of a lone analyst content to line up single jobs, push a button, get on with their evening and return in the morning to see the results.

But in the world of multi-person simulation teams in the aerospace, automotive and fast-moving consumer goods industries, with ever-decreasing product cycles, more variations in product lines and ever-increasing regulatory demands to be satisfied, the queue is clogged with hundreds of simulations, and the clock is ticking very loudly. The valiant workstation, chugging away on all processors, cannot keep up.

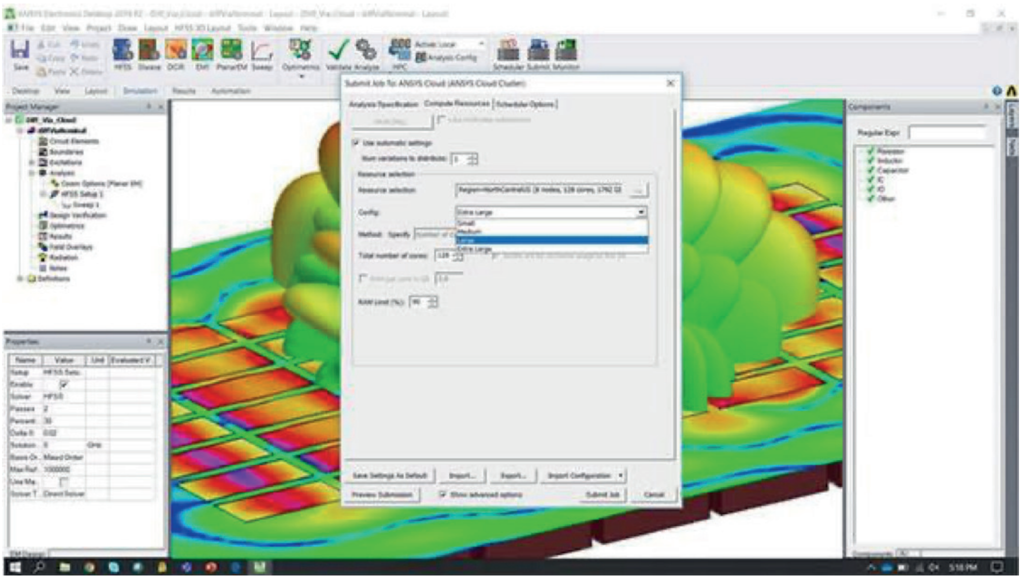

Submitting a solution to Ansys HPC cloud can be done through the Ansys interface. (Image courtesy of Ansys.)

WHAT COMPUTER HARDWARE IS BEST FOR SIMULATION?

The question of what kind of computer you should get for simulation used to have an easy answer: you got the fastest hardware you could afford. You demanded your manager get you a true workstation, not a PC like the other office workers. You had a lot of numbers to crunch, and the faster you could crunch them, the more simulations you could do. It was simple. While PCs spent most of the time waiting for commands from office workers, you were always waiting for your computer. Your manager may brag/complain of a spreadsheet so complex it takes a full minute to resolve the hundred formulas crammed into it, but you know you need to solve millions of simultaneous equations and that is going to take days—or weeks.

A top-of-the-line workstation can cost about ten thousand dollars and any of the big three computing companies—Dell, HP and Lenovo—will be happy to configure one for you. Prices start a little over $500 (Dell) for a machine that can solve textbook problems. For the real world, you'll need more.

The ultimate Intel/Windows-based workstation at the time of this writing may be the BOXX APEXX D4, configured especially for simulation and loaded with 10 CPU cores. This one starts at $8,290 but if you go wild with the configuration it will set you back over $50,000.

HARDWARE IN THE BIG LEAGUES OF HPC

HPC is a logical step up from a workstation for engineering groups doing simulation. (Picture courtesy of Ansys)

Should a workstation prove to be insufficient in terms of speed, or if your boss says you are not spending enough (we can dream), the next level of hardware speed comes with high performance computing, or HPC. Not only are all the computer components—such as the CPU, GPU, memory and storage—made industrial strength, but there are more of them. HPC requires high speed connections, parallel connections and a special operating system (usually Linux) to tie everything together.

Moving from workstation to HPC is a big jump in all respects: performance, size, cost and comfort. Your computer will have relocated from under your desk onto a rack. You won't have the Windows environment you were comfortable in. The computer itself is so large it will have to move to another room and be looked after by another person, someone more familiar with its Linux language and its care.

PERFORMANCE PAYOFF

Separation blues from not having a computer under your desk will soon be offset by the sheer power of HPC.

HPC capacity is measured in number of cores. While a robust workstation may have 8 CPU cores, and a server can house many more, a single HPC cluster cabinet can house a couple thousand cores.

An HPC unit will resemble servers—almost always rack mounted, and with the rack sometimes enclosed in a cabinet. A single rack or cabinet is called an HPC cluster. Because you can connect clusters to each other, HPC is infinitely scalable.

Expect to pay anywhere from $50,000 for a low-end cluster (an HPC cluster appliance) to $10 million for an HPC data center—a large room with rack after rack of HPC hardware. Building the room, which needs its own power and cooling, would be an extra cost.

THE RISE OF THE LOW-COST COMPUTING NODE

One approach that big companies are using to minimize unused and expensive resources is by making use of centralized computing with decentralized, low-cost computers on engineers' desks. This is hardly a new concept. It harkens back to the monitor/mainframe configuration that first brought computing to engineers—a concept upset by the advent of the personal computer, but now set to make a comeback.

At its modern extreme, a Chromebook will be the engineers' data entry and viewing, with all processing done somewhere else: on-premise, off-premise or way off-premise.

Making a virtual desktop on your local node, while all the software on your desktop lives somewhere else, is called desktop virtualization. It has been created and promoted by companies such as NVIDIA, Dell, VMware, Citrix and others. The idea is that racks of servers are connected to low cost computing nodes connected via Internet.

"Remote visualization and virtual desktop infrastructure (VDI) become a topic once end-users want to take advantage of centralized, shared and remote compute resources," says Slagter. "Since we want to give users the highest performance and most reliable options, we officially certify and support a range of remote display tools and VDI configurations."

Dell's Wyse "thin clients" or "zero clients" exist only to send keystrokes and mouse movements to the real computer and receive pixels in return. They look more like docking stations than computers; essentially, that's what they are, with a monitor, mouse and keyboard attached to make it all work. A thin client starts at a little less than $400, but when you add up all the components, it is hardly a compelling answer to a PC.

PC OR WORKSTATION, WHAT'S THE DIFFERENCE?

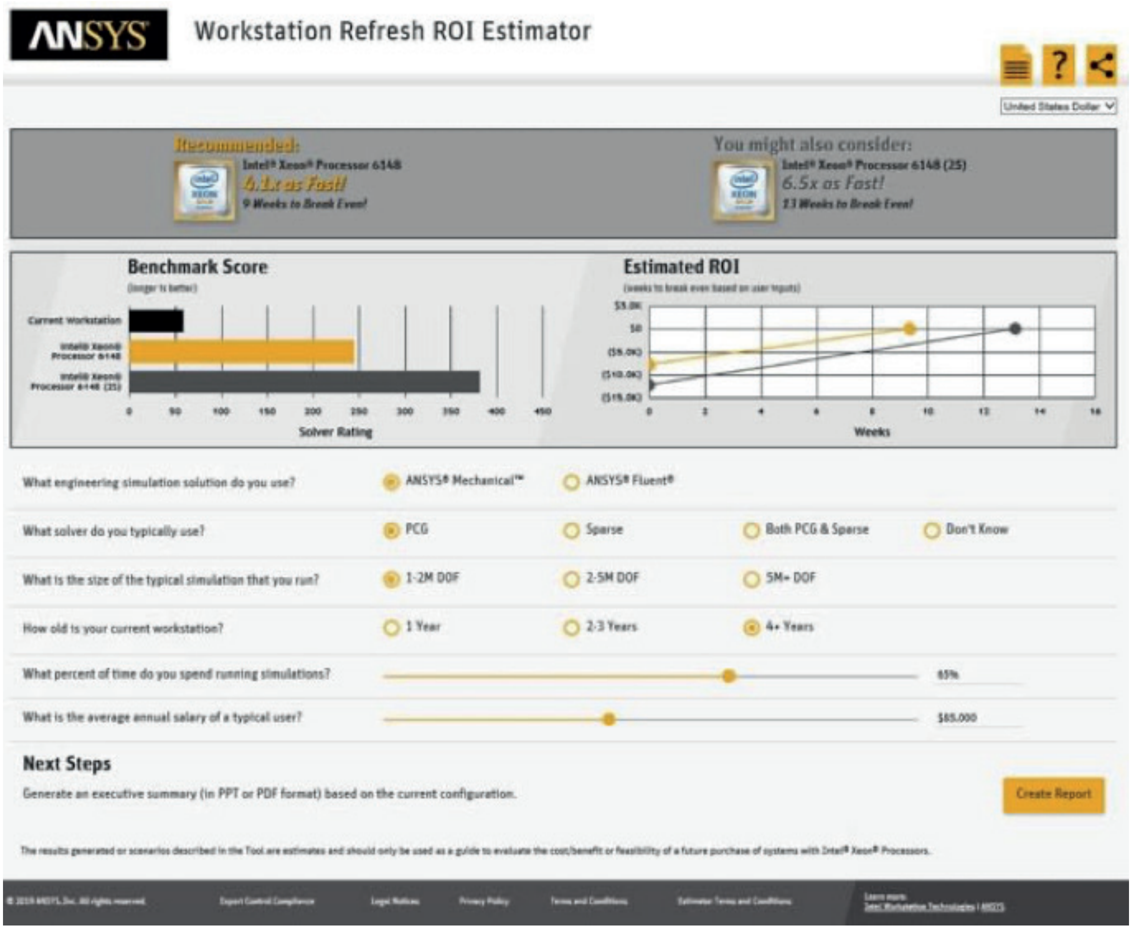

The Ansys ROI calculator will determine how long before new fast workstation will pay itself. (Picture courtesy of Ansys)

On seeing the tremendous disparity between computers costing well under a thousand dollars and workstations that start well above that, you might be thinking, "What's the difference?"

The price of inexpensive, basic PCs will certainly tempt the budget-minded. Or, you may be eyeing a sleek and svelte laptop like the Microsoft Surface Pro or a MacBook. However, unless you're solving trivial simulation problems, the PC or Mac simply won't do. (If you are doing simulation on the cloud, it's possible. More on that below.)

Simulation demands a workstation. It will lurk under your desk, lights glowing like a mad dog that will tear into any simulation you give it, spitting out results in no time—simulations that would have choked a PC.

A workstation differs from a PC mostly in the differences seen in each of the major constituent parts: a faster processor, more memory and more storage. Let's look at each class of parts in more detail.

PROCESSORS AND MEMORY

Taking advantage of the latest processor family is vital for CPU-starved simulation users. Yet, in a recent survey, Ansys found one out of six users are chugging along with workstations that are more than three years old.

A workstation-class processor usually has a faster clock speed and more cores than its PC counterpart. What may not be known is how many errors a PC will make.

If you think computers don't make mistakes, you're not alone. The history of computers making mistakes is almost as old as computers themselves. The "bug" in computer terminology famously comes from a scorched moth that caused a short circuit between two relays in the Harvard Mark II computer in 1947.

Another threat comes from cosmic rays which penetrate the atmosphere and are known to flip bits in computer memory. If you know computers—as Intel does—you suggest error correcting code (ECC) in your computer's memory.

A computer uses memory to make calculations; if that memory misfires, the calculation is wrong. The odds of a memory error are small for a single calculation, and that might be acceptable for a miscalculation that blanks a pixel on your screen and may not happen in a hundred years on a spreadsheet. However, if OS code is incorrectly compiled, it could be what keeps crashing a computer.

Does this sound like Chicken Little saying the sky is falling? Aren't we protected from most cosmic rays by Earth's atmosphere?

Those who have reasons to worry take notice of cosmic rays. NASA, whose projects get hammered by cosmic rays—and whose Hubble Space telescope problems may have been caused by cosmic rays—use ECC memory. So do avionics manufacturers whose work operates in the upper atmosphere.

If you assume the worst conditions, the effect of cosmic rays may occur in some form or another, probably undetected—a glitch here and there, or another Windows reboot every three days, according to a calculation by Berke Durak. Durak has a doctorate in theoretical computer science, and his calculation is cited by Intel because ECC memory fixes 99.988% of memory errors.

CPUS AND GPUS

The first computers had a single CPU, or central processing unit—essentially the brains of the computer, with everything else serving other purposes such as memory, storage, power, connections, monitor, keyboard and more.

A floating point processor came along a little later to assist with calculations. In 1992, the term GPU was first used by Sony for what was inside its Play Station and the GPU, or graphical processing unit, was born. In 1999, chip maker NVIDIA ran with the idea of a programmable GPU that could be called on to assist with all sorts of calculations, not just those for games, and creating the general-purpose GPU.

Before the GPU, games were as flat as Pac-man. After the introduction of the GPU, we saw Myst, with jaw-dropping graphics all the way up to Grand Theft Auto—with mind-boggling violence, but in full 3D.

Few gamers had a big enough allowance to buy high priced gaming cards with GPUs, so making GPUs available for professional graphics was a gold mine for NVIDIA. Now architects can dazzle their clients with beautifully rendered, gleaming skyscrapers and dream houses. More recently, GPUs are useful for deep learning, which can employ GPUs' massively parallel architecture. While interactive CAD never really benefitted from GPUs, simulation programs can be written to take advantage of GPUs, with computations happening simultaneously, instead of sequentially.

But as NVIDIA became a powerhouse, CPU manufacturers—most notably Intel—began defending their processors. It became a fight, when each type of processor actually benefits from the other.

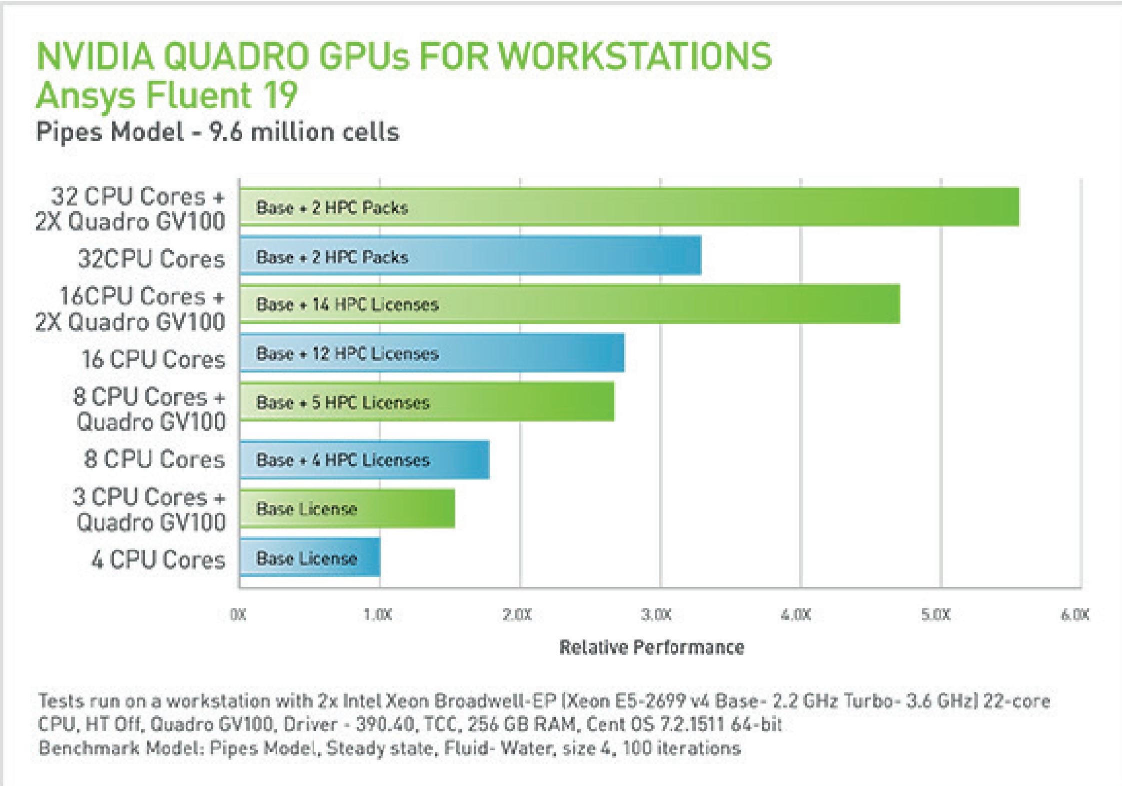

The effect of adding GPUs to a workstation dramatically increases performance for a 10 million cell CFD model, in this test by NVIDIA. (Picture courtesy of NVIDIA.)

A GPU cannot exist without a CPU; a CPU can reduce its workload with a GPU. An engineering workstation may have both fast, multiple core CPUs from Intel or AMD, in addition to GPUs, from NVIDIA or AMD.

Should your computer not have GPUs, the fastest way to increase performance is to add a graphics board with GPUs. Look for a GPU board with at least 5 GB of memory.

Not all GPUs are created equal. For Ansys software, certain NVIDIA models are favored for their ability to handle double precision calculations.

A CPU does a lot more than a GPU; therefore, the architecture of a GPU is relatively simple. A GPU is built almost entirely of pipes through which calculations flow, which leads to the description of a GPU as a graphics pipeline.

DRIVES

Simulation problems solve matrices that are too large to fit in RAM, so they are solved in parts that are constantly swapped in and out of storage—a process known as in-core and out-of-core solution.

Storage was traditionally the realm of a spinning hard drive. Even a 7200 RPM hard drive pales in comparison to the modern solid state hard drive (SSD), however, which has no moving parts and therefore theoretically operates with almost the same light-speed reaction time as RAM itself. This makes the connection between the SSD, RAM and the GPU and CPU cores critical—and also the source of bottlenecks.

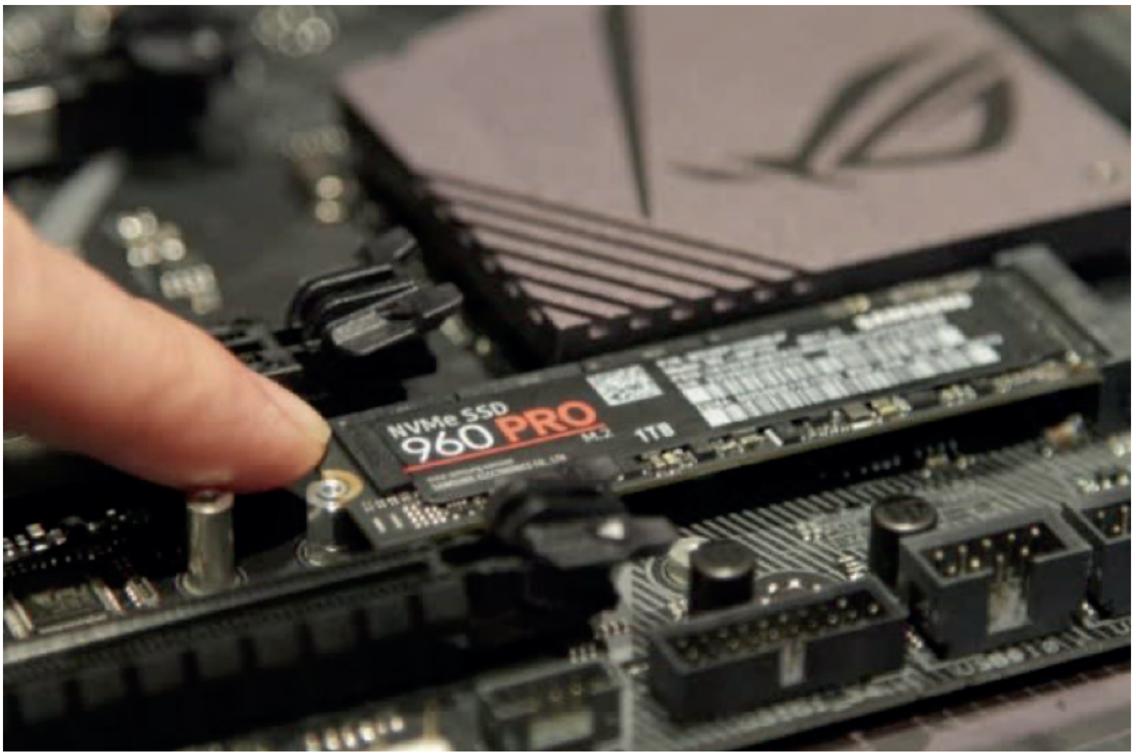

Not looking like a hard drive at all, an M.2 form factor SSD more closely resembles a memory board. Shown here being installed in a workstation, it is a great fit in laptops and mobile workstations, due to its diminutive size. (Picture courtesy of GroovyPost.com.)

SSDs typically have the same form factors as the hard drives they replaced, so that workstations could easily be upgraded. SSDs rely on SATA connections, which can't help but cause a bottleneck because calculations happen faster than the maximum 600 MB/s that the SATA III specifications allow.

Connect right to the motherboard using a PCIe connection is better. An SSD with a PCIe connection looks like a graphics board and takes up a slot in a workstation. Of course, this is not an option in a mobile workstation, which has no slots. A PCIe connection can have throughput of up to 4,000 MB/s—more than 6x that of SATA III.

The best solution appears to be SSDs that look more like a RAM module with an M.2 form factor, and which connect using an NVMe interface. Unlike interfaces that came before which were designed for mechanical drives with a single processing queue, NVMe is truly modern with thousands of processing queues. Throughput can be a staggering 32GB/s – more that 50x that of SATA III and 8x that of a PCIe.

A 1TB M.2 SSD can be added to a Lenovo P330 workstation for $435 at the time of this writing.

However, the large capacity SSDs common in workstations are still smaller than their hard drive counterparts.

THE HPC CLUSTER APPLIANCE—PLUG AND PLAY?

An HPC network can vary in size from a room full of computers shown earlier to the form factor of a single server, as shown in this HPE cluster "appliance" configuration. (Image courtesy of Ansys.)

An engineering group of four or five full-time analysts deserves an upgrade to HPC. They should arm themselves with an idea of the cost before they approach management asking for it. The good news is that in recent years, the price of HPC has tracked with the price of ordinary computing and gone down at a similar rate. And a new mini version of HPC, called an HPC cluster "appliance," has come onto the scene.

The HPC cluster appliance is no bigger than a rack mounted, conventional server—about the size of two stacked pizza boxes—and may be an affordable transition for the consultant or consulting firm. The benefit of a cluster appliance is that all necessary parts are included in the box. A total HPC configuration includes "compute nodes" and a "head node" as the physical high-speed interconnects between the two, and the management software, which is usually UNIX-based. A login node is included in the head node.

One such HPC appliance is offered for lease by HPE.

If you don't want to install and manage the HPC appliance yourself, the company TotalCAE offers to wheel it in, plug it in, get it up and running and maintain it for you.

Although several companies are willing to sell you their version of an HPC appliance, we found few were willing to openly share their prices. It was always "depends on the particular configuration" and "please contact sales." We understand an HPC appliance is not like buying a fridge at Home Depot, but having to speak to a human for pricing is anything but modern.

An entry-level HPC appliance, the Essential by TotaLinux looks like a normal Windows-based workstation and sells for about $11,000 USD, not including software. (Picture courtesy of TotaLinux.)

One vendor who was public about prices was France's TotaLinux, with HPC appliances starting from 9,990 € (about $11,000 USD) not including software for their 32-core Essential to the 96-core Ultimate for 44,990 € (over $49,600 USD).

Be warned that although its creators may want to make stepping up to an HPC appliance sound like a plug-and-play, DIY affair, it will be anything but that for an engineer used to a Windows based workstation. HPC requires a professional installation, and will involve considerable setup, tweaking and training. If you have a small group of engineers currently doing simulation, you can count on at least one of them learning LINUX and being unproductive for a week or two during the set-up period.

Creating and maintaining an HPC data center will make sense if your simulation needs are massive and regular. For smaller operations, an HPC appliance configuration may suffice. (Picture courtesy of Ansys.)

HPC APPLIANCE CONFIGURATIONS FOR ANSYS

BASED ON INTEL® XEON® GOLD 6242

	64-core server	96-core cluster
		 Non-contractual picture
CAPEX	**~ $ 21 500**	**~ $ 48 500**
OPEX (price per month over 3 years)	**~ $ 430**	**~ $ 1 500**

“We have partnered with hardware and service providers, including Dell EMC and HPE, to release what we call ‘plug and simulate’ clusters,” explained Slagter. “These are preconfigured with Ansys simulation and job management software. Think of them as ‘cluster appliances’ because they are basically turnkey appliances that reduce the time and cost of putting an HPC system together—finding and buying the hardware, configuring, testing, etc. A cluster appliance is very easy to acquire.”

“And a cluster appliance is easy to maintain,” he added. “It can be maintained by internal IT, if a customer has one, or it can be externally managed by one of our partners or system integrators that are part of our HPC system.”

“HPC is not easy to get, or easy to use,” admitted Slagter. “But with our HPC partnerships, we are trying to take away the barriers to HPC adoption.”

For this reason, getting an HPC appliance on an account that includes installation, maintenance and initial training will be a good idea.

ENTRY-LEVEL HPC CONFIGURATIONS

For those who have ten or more engineers doing simulation full time, the next level up from an HPC appliance—a mid-level HPC installation—may make sense.

Unlike the world of commodity PCs and workstations, where products compete on prices and prices are readily available, the lofty world of high-end computing is far more secretive. It’s not about selling you a product as much as it is about providing a service, with a product (the HPC system itself) that is custom made to your needs.

For the modern engineer or manager, who is used to shopping online, this can easily lead to the belief that HPC is for big companies that can absorb a cost-is-no-object system—no doubt a multimillion-dollar room full of servers. Who wants to hear, “If you have to ask, you can’t afford it” from a haughty salesperson?

An entry-level HPC configuration with 560 nodes for $150,000 from ACT Systems. (Image courtesy of ACT Systems.)

But we found one company, Advanced Clustering Technologies, that bundled its HPC offerings into entry-level, intermediate and advanced configurations—and made the cost no secret. While that may give you sticker shock if you are accustomed to buying workstations at a hundredth of that cost, consider $500,000 to be the rough cost of four engineers in the U.S. doing defeaturing for one year.

Advanced Clustering Technologies, or ACT, is based in Kansas City, Missouri. ACT offers its systems in neat $100,000, $150,000, $250,000 and $500,000 packages. Unlike the entry-level HPC appliance, the individual head and compute nodes are separate units and take up multiple rack positions.

Assuming $150,000 is an entry-level rack-mounted HPC configuration, ACT would provide:

- 14 server “blades,” each with 2 Intel Xeon Gold 6230 processors, each with 20 cores, for a total of 560 cores
- Memory, storage and additional hardware
- HPC network software

Installation will require a professional, as the $150,000 system is a 1,000 lb. behemoth that uses 10kW of power, requires three 220V, 30 Amp circuits, and generates 36,000 BTU of heat.

SUPER COMPUTERS

At the extreme high end are supercomputers. The most well-known name in supercomputers, Cray, is still around but has been bought by Hewlett Packard Enterprise, or HPE. Other supercomputers are offered by IBM (Watson), plus Google, Fujitsu and a few others. But we will limit our discussion to workstations and HPC for this article.

RENT HPC

You can pay as much for a workstation as you can for a good used truck. But your Dodge Ram with full crew cab, full bed and dual rear wheels that you bought so you could haul your living room set home or your boat to the dock sits in your driveway most of the time.

Buying big for a one-time or occasional use doesn’t sit well with family members or company accountants, both having competing needs for capital expenditure. Unless the super workstation will be used all the time, it may be a better idea to rent one for your one-in-a-while need. Also keep in mind that all hardware has to be maintained, software updated, backups performed and aging hardware replaced—all labor that is attached to the capital expenditure.

You will be delighted to find out that the fastest computers are available for rental. Several companies have made HPC resources available on demand and charge only for the time you need. Instead of an enormous, gut-wrenching capital expenditure, HPC can become a monthly operational expense.

HPC appliances can be rented. For what HPE has to offer, see HPC Starter Kit for Ansys Environments and for other Ansys partners for rentable HPC appliances, go here.

HOW MANY CORES DO I NEED?

Early in the days of PCs, you could tell the performance of a computer simply by the clock speed of its CPU. In 2001, IBM created the first multicore processor and made this more complicated. Now there are multiple CPUs—called cores—on a single multiprocessor. The clock speed of any one core can be lower than a single-core processor, but combining cores more than makes up for it.

While thousands of cores are theoretically possible on a CPU, a recent survey by Ansys found the average workstation for simulation has dual processors with 12 cores total.

“If you buy a workstation for simulation, it should have at least 16 cores, i.e., dual 8-core processors,” says Slagter.

CLOUD COMPUTING

Some of our biggest tech companies, such as Google, Amazon and Microsoft, need computers on a global scale. This has led to the birth of cloud computing, a revolution in computing that has become part of everyday, first-world life.

Although the term “cloud computing” is said to have been coined by Compaq in 1996, it was Amazon that realized its excess computing resources could be made available for others. In 2002, Amazon Web Services (AWS) was born. The idea that a bookseller turned full-fledged shopping service would successfully rent out its excess computing resources and end up transforming computer use was just as much as surprise to Amazon as it was to the rest of the computing industry. By 2015, AWS was making Amazon almost $8 billion a year.

Being scooped by a shopping service must have been embarrassing for the existing tech giants, so they scrambled to join what consumers and businesses proved they wanted. While the market is dominated by Amazon, Microsoft has moved up to second place and Google—who may have more computers than anyone on Earth—is in third place. On the other side of the globe, China’s Ali Baba also offers cloud services.

GET A PRO

While the concept of sharing a large collection of computers is a basic one, the typical engineer will likely turn away from public cloud providers if they have to develop their simulation environment on their cloud infrastructure themselves. Building up an enterprise-grade simulation environment with HPC, job scheduling, graphics workstations and storage in the public cloud can easily take many months, as it is beyond their expertise or training.

A middle layer of services has sprung up to make use of the cloud more accessible to the public, which includes software vendors who interact with cloud service providers.

“The engineer’s job description defines productivity as figuring out product design problems, not cloud setup problems,” says Slagter. “This is why we have developed Ansys Cloud.”

GO FOR SPEED

For most of our readers—such those that are consulting or in small consulting firms or simulation groups, and those who are bumping up against the ceilings imposed by personal workstations, to whom high performance computing is tempting but unaffordable—the logical first step would be to avail themselves of simulation-specific HPC on an as-needed basis, through a companies like Ansys.

If you are using Ansys solutions already, you can access the Ansys cloud without leaving the Ansys application you are in. Besides pay-per-use usage, Ansys elastic licensing can now also use customer’s on-premise lease or perpetual licenses.

“Currently Ansys Cloud supports Mechanical, Fluent and Electronics Desktop products but in Q1 of 2020 we will expand the support to CFX and LS-Dyna,” says Slagter.

Simulation is very demanding of hardware; for that reason, the fastest computers are the best. For some engineers, simulation may best be done on a high-end workstation. Circumstances may dictate a more powerful server or an HPC appliance, or even a full HPC configuration with a rack or multiple racks, depending on the amount of simulation that needs to be done.

For an increasing number of engineers, as well as the small engineering firms dedicated to analysis, those who feel constrained by their personal workstations, or for organizations allergic to large capital expenditure, cloud HPC will appear to be the perfect step up. When you add the pluses of contracting an HPC service, such as always having access to the latest, fastest hardware, and not having to divert precious engineering resources to IT functions like maintaining hardware and software systems, it can be a no-brainer. They won’t have to worry about running out of capacity, as cloud HPC centers are able to scale up as needed. We’ve never heard of a simulation running out of room on a cloud HPC center.

Most importantly, you can have the super-fast throughput, with your results coming back in minutes instead of overnight or days later—and you can get started in as little time it takes to open an account.

HPC USERS: H IS FOR HAPPY

“The Ansys Cloud service built into Ansys Mechanical provides intuitive, easy access to HPC directly from the application. For large, high-fidelity models, Ansys Cloud reduced our solution times by 5 and 6 and cut the entire simulation workflow by half,” says Marcos Blanco, mechanical Simulation engineer at Lear Corporation

“High-efficiency equipment is critical for improving plant performance in the oil and gas industry. Ansys Cloud enables Hytech Ingeniería to calculate large and complicated geometries within hours, instead of days or weeks—resulting in significant time savings,” says Luis Baikauskas, process engineer for Hytech Ingeniería.

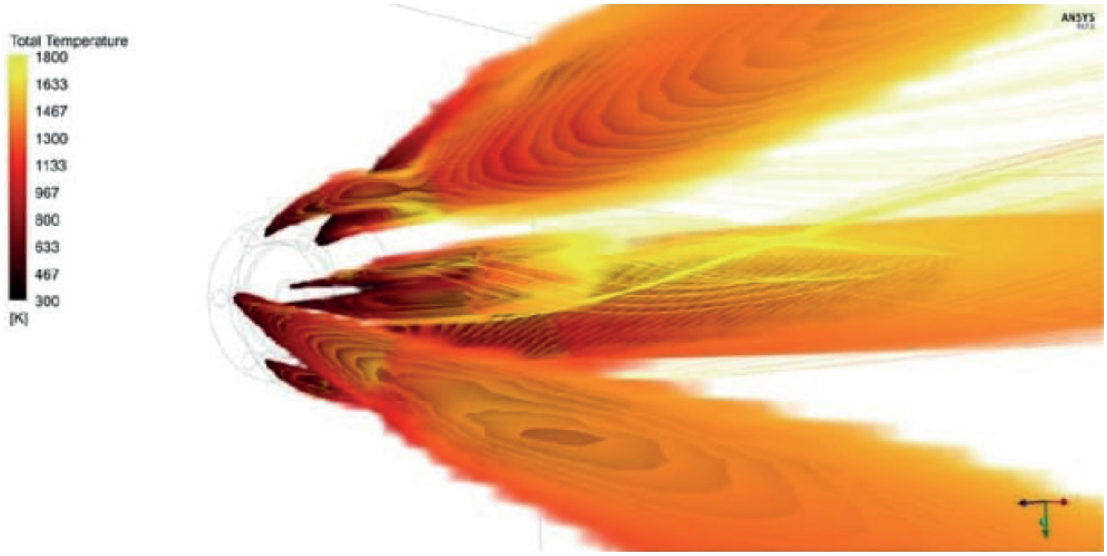

RJM CleanAir Gas Burner combustion simulation. RJM found simulation that took a week could be done in a day using HPC. (Picture courtesy of Ansys.)

RJM International are makers of industrial burners for power generators and other large combustion plants, such as those found in refineries and steelworks. RJM was doing combustion simulation, and the scale and complexity of the burners required simulation that took a week to solve. Moving the simulation to HPC cut the time by 86 percent, with simulations now being done daily.

Modine Manufacturing, a maker of heating systems, found its 10-core, 64 GB of memory, personal workstations would take three months to solve a transient heat transfer solution. The company could not tie up its workstations for weeks, much less months, so it had to look for a faster solution. Using Rescale’s HPC, the company was able to reduce the solution time to days. An additional bonus to Modine was in saving much of the time it took to build physical prototypes for testing because it was able to simulate instead.

Rolls-Royce is best known for its legendary automobiles, but you’re more likely to be transported by the company’s jets engines than its internal combustion engines. The company is the third biggest manufacturer of turbofans (behind GE and Pratt & Whitney). The aviation industry was the first to use FEA, and continues to make big demands on simulation software.

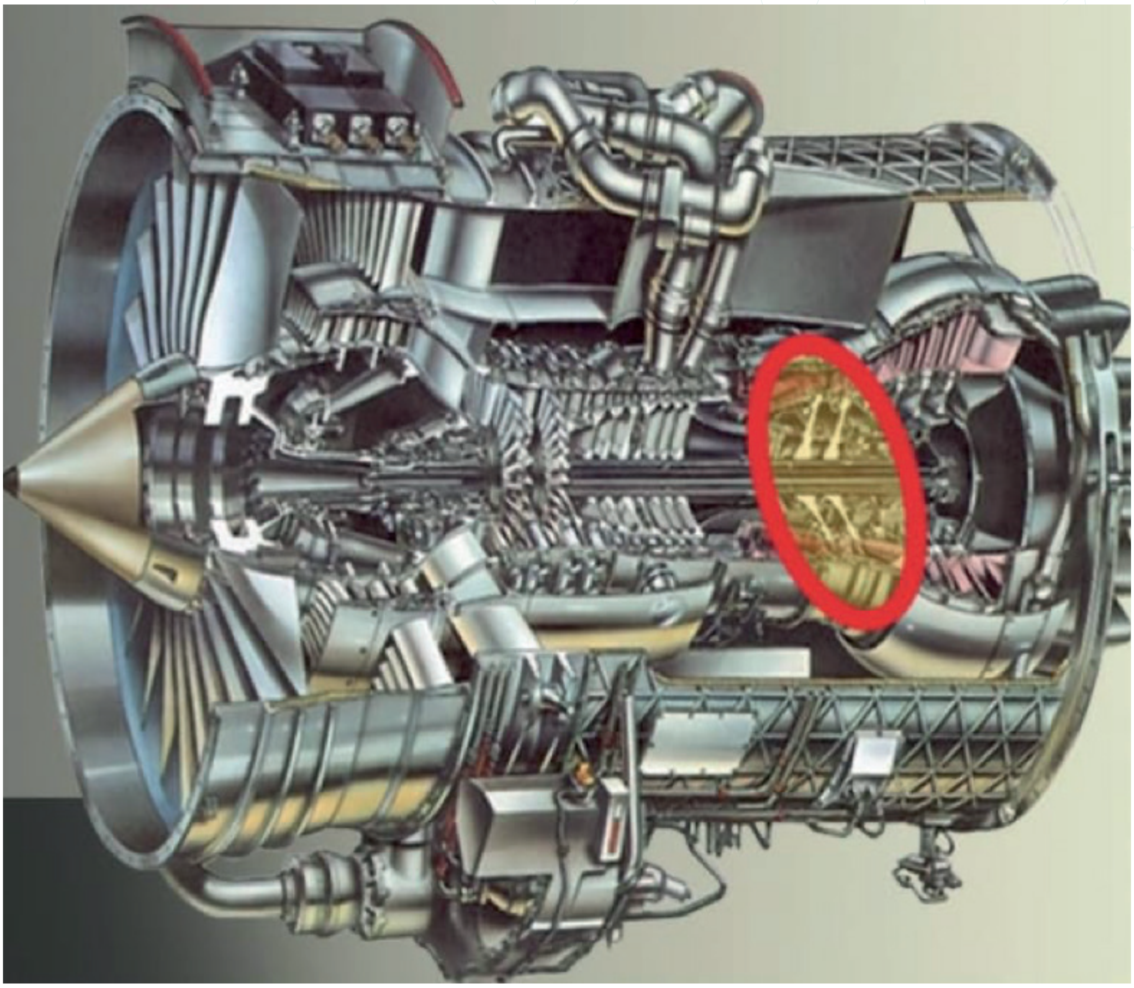

Rolls-Royce turbofan showing the interstage cavity where the heat produced by the fluid flow is used as a boundary condition for the structural members. (Image copyright Rolls Royce The Jet Engine.)

Rolls-Royce turbofan showing the interstage cavity where the heat produced by the fluid flow is used as a boundary condition for the structural members. (Image copyright Rolls Royce The Jet Engine.)

Rolls-Royce, no stranger to HPC, found itself maxed out of internal HPC resources when performing multiple multiphysics iterations of fluid flow coupled with structural simulation on its latest turbofan design. The company had to find HPC resources outside its walls, settling on HPC service provider CPU 24/7 GmbH. The simulation utilized 32 cores for the CFD while only needing one core for the structural simulation. It was a move that led to a five-times reduction in calculation time, according to Ansys.

Other benefits included avoiding costly design revisions during the production of the prototype and successfully developing a pathway for the elimination of thermal sensors in its engines.

LOOKING FORWARD

Engineering simulation appears to be bottlenecked by the computer hardware it runs on. For many engineers in a consulting practice or in small and medium sized businesses who do an occasional simulation, a fast, top-of-the-line workstation may be enough. Care must be taken so the workstations are kept up to date. However, for those engineers with a steady diet of simulation and those with large and complex problems, the next step will be high performance computing, or HPC.

With HPC, small businesses and entrepreneurs will be able to tap into the potential power and unlimited capacity of clusters of cores just like big businesses do, as HPC is available on-premise or in the cloud from software as a service (SaaS) vendors. As more SaaS services proliferate, it will become accessible to all.

Ansys is so convinced that HPC will help engineers who do simulation that it will run your Ansys CFX, Fluent, HFSS, Maxwell3D or Mechanical simulation on an HPC configuration so you can compare its times to what you get on your laptop or desktop workstation. And if engineers want to try HPC on the cloud themselves, they can go to www.ansys.com/cloud-trial.

Ansys invites you to see the difference with HPC for yourself with its free benchmark program, which uses "expert" configurations of high-end, simulation-capable Windows 10 workstations—one with Intel Xeon Gold 6148 Processor with 16 cores, 192 GB of memory and 2 TB of storage that cost $8,000, another with Intel Xeon Gold 6148 (2S) Processor with 32 cores in use, 384 GB of memory and 2 TB of storage for $12,000 and a third with a cluster of 128 cores.

"We will run the model and tell you how long it takes so you can compare the turnaround time to your workstations," says Wim Slagter.

Ansys

Learn more about high performance computing at Ansys.

engineering.com