Technical white paper

Beyond the desktop

Leveraging high-performance computing to turbocharge engineering productivity

Contents

Changing the landscape of computer-aided engineering

Over the past 10 years, computer-aided engineering (CAE), and more generally virtual prototyping, has been increasingly recognized in product design. Simulation-driven prototyping fuels innovation in the manufacturing industry and helps users explore real-world behaviors of products. CAE helps enterprises innovate to:

- Compete
- Cope with increasingly demanding regulations
- Improve product capabilities and quality
- Reduce costs and development times

Time is business critical. Time is the driving factor behind how we target business goals, measure profits, and track expenses. Time-to-market impacts product success and market acceptance. Productivity improvements boost results. Technology has only made us more intolerant of delays. Your staff doesn't tolerate delays and downtime in its personal use of mobile devices, or in accessing office applications for work—an attitude enabled by the pervasive use of networking, centralized servers, and storage.

However, for many businesses, engineers accept slow performance of key engineering applications thanks to out-of-date workstations. Model response times and the scale of analysis are constrained by the performance, memory, and storage limits of a single system. A shared, multisystem high-performance computing (HPC) resource provides much faster response times for simulations, supports larger model sizes, and supports engineering collaboration across the organization. HPC users report faster time-to-market, improved productivity, and deeper insight into product designs. Tapping HPC fosters success beating the competition, meeting the demands of increased regulation, and tightening staffing levels.

Fortune 500 companies, major universities, and researchers use HPC. In the past 15 years, the HPC market has grown from $4.8 billion to over $12.5 billion annually,[1] reflecting the increased adoption of HPC technology across the world. As reported by the Information Technology and Innovation Foundation, "high-performance computing has become a critical enabler of innovation, new product design and development, and product testing and validation in virtually all manufacturing companies, meaning HPC is helping manufacturers both cut costs and create new revenues."[2]

Small companies have been slower than larger firms to adopt HPC and advanced modeling, with a willingness to "get by" with older paradigms and ignore the hidden burden on productivity and time to market. Concerns about the cost of deployment and lack of in-house expertise hold back these firms from stepping up to HPC. While there is an upfront cost in capital investment and a learning curve to overcome for those who are new HPC users, the cost for an entry HPC cluster is lower than many companies may realize. Specific tools and services make stepping up to HPC easier than ever.

Accessible HPC

Advances in hardware and software dramatically lowered the barriers to supercomputing in the early 2000s. Use of industry-standard technologies, such as Intel® processors, Ethernet networking, and open source software allowed HPC to move beyond a community of well-funded universities and government labs to a global user base that crosses industries, education, and research. As the technology becomes more affordable and pervasive, tools and applications were commercialized—with simpler user interfaces and access to professional support services.

This is especially seen in CAE. Aerospace and nuclear science were early adopters of CAE, as the use of physical prototypes and testing were infeasible, dangerous and/or cost prohibitive. Similarly, oil and gas exploration benefited from modeling for design of drilling equipment, as well as seismic data processing. The successful use of simulation gained broader acceptance and drove the commercialization of in-house and academic tools. Today, CAE is used across manufacturing.

An entry-level HPC system is able to deliver powerful performance in a very small form factor that can be easily integrated with existing enterprise servers. For example, a four-node HPE Apollo 2000 Gen10 System can fit into standard IT racks and only takes 2U of space. With the most current multicore processors, such a system could have 160 cores, 1.5 TB of memory per tray for a maximum of 6 TB, and minimal demands on your power and cooling infrastructure.

[1] Worldwide HPC Server Forecast, 2016–2020, IDC, June 2016
[2] The Vital Importance of High-Performance Computing to U.S. Competitiveness, April 2016

Systems can be configured with a switch for a private network to minimize the impact on the enterprise network (and optimize HPC performance). The system operates independently of the enterprise IT resources—with its own private network. The resource is connected to the IT network via one node on the cluster.

The HPC system can also be easily (and remotely) accessed by IT or engineering for management, performance, and maintenance of the system. For the engineer, the HPC resource can be tapped via the application's interface on the desktop. Engineers can submit jobs and the batch scheduler will allocate resources within the system based on the engineer's specifications. Although most HPC systems use Linux®, Windows® is also a good choice for smaller sites unfamiliar with Linux.

System vendors and their value-added resellers provide services to assist with rapid deployment. Hewlett Packard Enterprise, for example, offers HPC configuration tools to assist in deploying a tested, integrated cluster. These tools are complemented by reference architectures (RA) for specific applications. RAs include entry-level designs and provide advice on processor selection, memory size, and storage selection. The combination of configuration tools and RAs reduce risk, along with the time to deploy and include options to allow users to customize based on specific site requirements.

Hewlett Packard Enterprise and its resellers can prepare the system prior to shipment, including the installation of the ISV software. They can also assist with on-site installation. HPE resellers can offer application-specific support and training, and provide ongoing management and maintenance.

Cloud computing has also opened access to high-performance computing. Cloud service providers offer the ability to specify node characteristics and define a virtual HPC cluster network. For some, HPC on the cloud can provide a test bed for experimentation with HPC. Applications vary on how well-suited they are for a public cloud. Concerns can include the size of the data over the network, security, and availability of commercial CAE applications at the site. However, the growth of interest in cloud computing is spurring improvements.

Benefits of HPC

Realistic simulation can push the need for capacity beyond a single workstation or server. Complex models place demands on storage and memory. Product designers may constrain the size and quantity of simulations to reduce wait times. Cost and management of multiple high-end workstations can often be improved with remote visualization from a central resource. Scaling out the workstation infrastructure by adding clusters will increase productivity, improve model fidelity, and simplify user access to the latest technologies (such as accelerators).

Faster simulation

Shorter time-to-market through reducing design and development time is often cited as one of the key benefits of simulation. An HPE partner, ANSYS, is the global leader in engineering simulation software. Using ANSYS® multiphysics simulation, Tecumseh engineers took only 18 months to design a new family of compressors, compared to the four to five years the team previously required to build and test dozens of prototypes. The shortened design time resulted in a 60 percent reduction in time to market.[3]

CAE accelerates the design process. The use of cluster resources accelerates simulation when compared to a single system or workstation. Hewlett Packard Enterprise, in collaboration with leading software developers such as ANSYS, regularly conducts benchmarks for CAE applications. For ANSYS Fluent, the speed up achieved through an HPC cluster has been demonstrated for over a decade. In a recent benchmarking test by Hewlett Packard Enterprise, simulation times were compared running a Fluent standard benchmark on a single node, up to a 72-node configuration using the Intel® Xeon® Gold 6148 Processor 2.40 Ghz 20 core processor with Turbo on.

[3] Breakthrough Energy Innovation Volume X, Issue 3, ANSYS Advantage, 2016

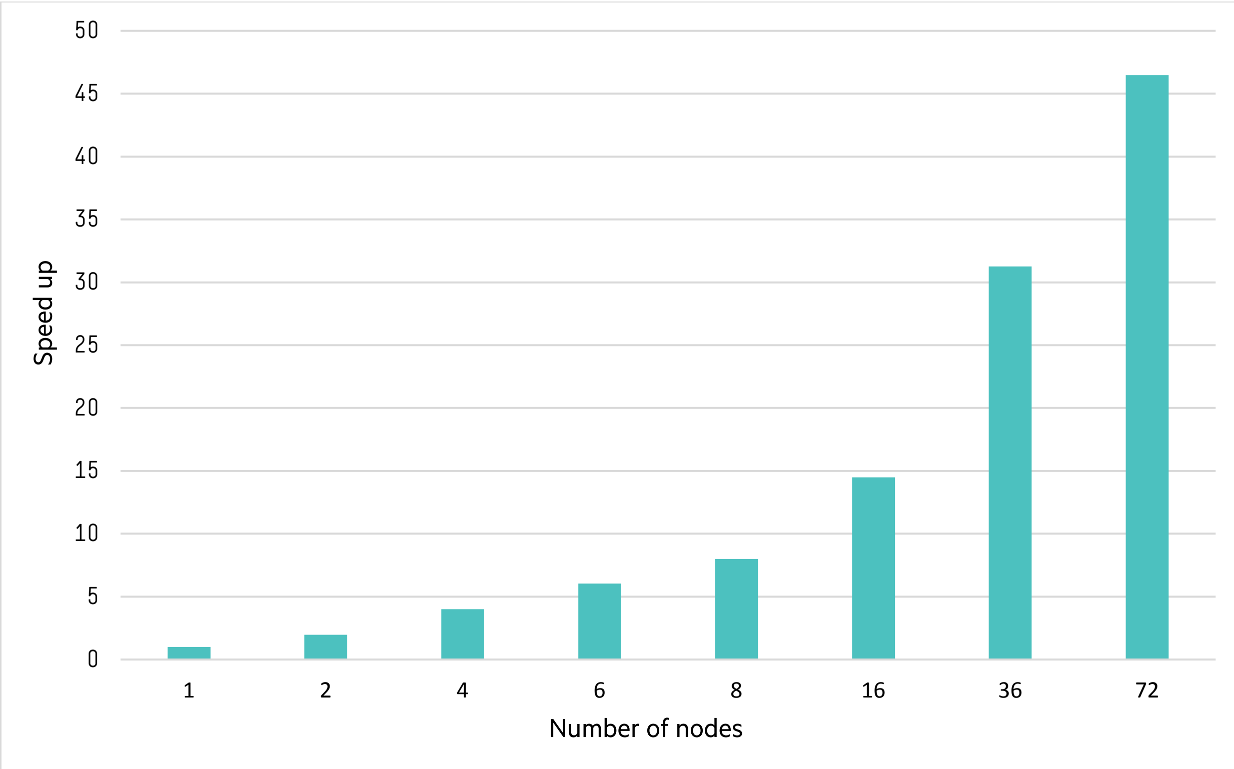

Figure 1. Speed up of ANSYS Fluent for small and medium benchmarks[4]

At four nodes, for example, you see a 4X speed up. In essence, you could run the benchmarks on a four-node cluster in 25% of the time it takes to run those same benchmarks on a single node (or workstation).

A related benefit from faster simulation is the ability to run more models, in the same time frame. For example, you could run 4X the simulations on the cluster. This allows engineers to examine more design variants, and explore and optimize designs faster. Either way, users will see a tremendous boost to productivity.

Larger simulation models and data sets

The size and performance of the largest computing systems increase every year, as the world's top researchers harness technology to address science's grand challenges. For many, simulation is accomplished on a more modest scale. However, the common experience reported is that analysts are always eager to tap into new technology and compute capacity to run larger models to gain a better understanding.

The evolution of simulation models reflects the evolution in computing capability. Reservoir models used in oil and gas exploration evolved from 2D to 3D models with more realistic geometry. Enhanced graphic capabilities allowed developers to visualize and manipulate designs. Increased processing power, memory size, and faster storage access allowed engineers to tackle larger and more complex problems, and move from simplistic models to "systems of systems." Advances in software also enable an interdisciplinary approach toward design integrating multiphysics, materials, and life sciences.

The size of finite element analysis (FEA) and computational fluid dynamics (CFD) models applied to product design can exceed millions of data elements. For many models, polygonal or polyhedral meshes are used to approximate geometric domains. Typically, finer meshes will generate more accurate models, but will also increase the size of the model and the computational time required.

[4] Test using ANSYS Fluent 19 standard benchmark Aircraft Wing 14 million cell model, Hewlett Packard Enterprise, March 2018.

Designers will also want to model behavior varying external conditions such as dust, humidity, shock, and vibration adding to model complexity and size. To deal with model size and complexity, one approach has been to analyze at the component level and integrate pieces at the final stages of development. This significantly impairs the predictive accuracy and realism of the model.

HPC allows users to tackle larger, more complex models. Recommendations from CAE software vendors like ANSYS provide guidance on the number of cores, memory, and disk, as data elements or degrees of freedom increase. Actual performance will depend on the model. For small systems, a general guideline is that a four-node HPC platform, such as the HPE Apollo 2000 System, should support models 4X than those on a single node.

Reduced physical testing

A primary driver of early simulation was the need to substitute physical tests where such tests would be risky, costly, or infeasible, such as nuclear science and aerospace. A reduction in physical prototyping produces savings in time and dollars. Use of virtual prototypes earlier in the design cycle yields savings in total design time. Virtual prototyping reduces (if not eliminates) the number of physical test models and operating costs for the test facilities. In aerospace, wind tunnels use has been replaced almost entirely by CFD. Crash simulation software allows engineers to improve vehicle safety and understand collision impacts, and refine designs before final physical tests.

Consolidation savings

A long-established benefit of clusters has been increased utilization of resources when compared to isolated, individual workstations and servers. Pooling of resources allows a team to leverage the capacity of the entire group. The shared resources can include specialty technologies such as accelerators, which may only be useful for specific applications and models. Application software licenses can be shared as well.

Savings are achievable by reducing the requirement for high-performance workstations. With access to an HPC resource, individual clients do not need to be maxed out in memory and storage. The HPC cluster can run resource-intensive simulations, while pre- and post-processing takes place on the client level. In addition, adding remote visualization capabilities will enable post-processing applications right in the cluster where the data is produced. Per seat software licensing is frequently lower when deployed on clusters. For example, Hewlett Packard Enterprise has worked with Red Hat® and SUSE to offer specially priced operating system subscriptions for HPC clusters. Middleware and applications for HPC environments are often available with volume discounts for cluster deployment.

HPC costs

HPC offers compelling improvements in engineering productivity, time-to-market, and product quality, which offset the investment in HPC technology. IDC is currently conducting a multi-year study of HPC and ROI with funding from the U.S. Department of Energy. In a 2016 update, $44 of profit is associated with $1 of investment in HPC.[5] Meanwhile, the cost of deploying HPC continues to fall due to increased system performance, simpler user interfaces and cluster management, and new usage models such as cloud computing.

Benefits of HPC

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

Various ANSYS HPC licensing options allow scalability to whatever the computational level simulation requires; from single-user to small user group options to enable entry-level parallel processing up to virtually unlimited parallel capacity. For large user groups, ANSYS facilitates multiple parallel processing simulations, highly scalable for the most challenging projects when needed. Apart from parallel computing, ANSYS also offers special product solutions for parametric computing.

[5] HPC update at ISC 16, IDC, June 2016

Initial investment cost

Entry into HPC can be supported with systems as small as two to four nodes. With core counts as high as 20 cores (with Intel Xeon Gold 6148 2.4 GHz processors), impactful parallelism is achievable even in small clusters. The key components within the cluster are the compute nodes, networking, and software.

As noted earlier, the HPE Apollo 2000 System includes up to four nodes within a 2U chassis that can be easily added into existing on-site racks. Sites anticipating expansion or desire to deploy systems as a free-standing resource may use a dedicated rack. The HPE Apollo 2000 System utilizes a 2U chassis with a shared power and cooling infrastructure providing power efficiency and cost optimization. The chassis can hold up to four two-processor nodes.

To exploit the power of today's systems fully, Hewlett Packard Enterprise supports all major commercial interconnect technologies giving customers a flexible choice to optimize price and performance. These technologies include Ethernet, InfiniBand®, and the Intel Omni-Path Architecture (OPA). In addition to the cluster switch, Gigabit Ethernet should be utilized for non-compute traffic within the cluster (supporting the administrative tasks for the systems and the out-of-band network) and to support the connection with the site's networking infrastructure.

The software stack for a cluster typically includes the operating system, cluster management, and a job scheduler. In addition, users may include performance tools and libraries. A message passing interface (MPI) is often included with applications, but can be deployed separately. The cost will depend upon the specific system configured to fit the user requirements.

Integration and installation

From a hardware perspective, assembly and physical deployment of a small cluster is straightforward. For example, the HPE Apollo 2000 System chassis holds the nodes and the chassis fits into a standard rack. A Gigabit Ethernet switch is connected to the Ethernet and management ports on each node, while the IB or OPA switch is connected to the IB or OPA enabled HPE FlexibleLOM. One of the nodes serves as the head node. Typically, that head node will be the portal into the HPC resource and can be connected to the enterprise network. Customers often opt for factory and/or on-site integration, especially for larger systems, and incorporate the software installation.

Software installation can be enabled with cluster management tools such as HPE Insight Cluster Management Utility (CMU). HPE Insight CMU is used by customers to provision, manage, and monitor Linux-based clusters. For Windows, Microsoft® HPC Pack provides well-documented tools for deployment and maintenance. However, an HPC environment uses tools and libraries that are specific to technical computing, creating a learning curve for those new to HPC. Hewlett Packard Enterprise and its resellers offer start-up services to help customers with this transition.

The impact on the IT infrastructure of adding an HPC resource should be minimal. The system is compact and can be installed into existing racks. The power consumption will depend upon utilization rates, but with only four nodes, the system should not impact the power or cooling requirements for the IT infrastructure.

Ongoing maintenance and support

As with any system acquisition, one needs to budget for ongoing hardware and software support. System vendors such as Hewlett Packard Enterprise can be a single point of contact for integrated systems, including operation systems and key tools such as the scheduler. Application support, however, needs to be delivered by the software vendor or its partner service providers.

Key considerations

A transition to HPC is more successful when sites have analyzed their current and future requirements, assessed their readiness, and taken necessary measures to prepare for deployment.

System design and specifications

Applications place different demands on memory, disk access, and core counts. Some applications may not generate the network traffic created by MPI-intensive jobs. If the site plans to support a mix of applications, a general purpose HPC cluster can be designed to support such a mix. Different nodes can be configured to support targeted applications. For example, some nodes can be configured with GPU accelerators or with larger memory. A robust scheduler would be able to align those nodes to the jobs that can best exploit those capabilities.

Starter cluster recommendation

A typical cluster configuration allowing expansion up to four nodes includes the following baseline:

Table 1. Recommended small cluster configuration

Cluster baseline	Designed for up to four nodes in one HPE Apollo r2x00 chassis
Head node	Either 1 HPE ProLiant DL360 Gen10 head node (external) or a single HPE ProLiant XL170r (within the HPE Apollo 2000 chassis)
Compute nodes	Two to four HPE ProLiant XL1x0r nodes One HPE Apollo 2000 holds up to four HPE ProLiant XL170r nodes or two HPE ProLiant XL190r nodes. Customers may also mix and match based on their application needs. Up to 40 cores per compute node, Intel Xeon Gold 6148 20 core 2.4 GHz processors recommended
Administration network (out-of-band and console)	One Aruba 2920 Switch Series with 10GbE uplink
Interconnect	10GbE, InfiniBand, or Omni-Path Architecture
HPE Apollo r2000 chassis	One HPE Apollo r2x00 chassis, each with 4 fans, and 2 x 1400W power supply
Operating system	64-bit Linux and Windows
Cluster management	HPE Insight Cluster Management Utility
MPI	MPI software is included in ANSYS installation package
Job management	ANSYS Remote Solve Manager (RSM) and supported job schedulers: PBS Pro (Linux) Torque with Moab (Linux) UGE/SGE (Linux) Windows HPC 2012 R2 Platform LSF (Linux)
Remote visualization	NICE Desktop Cloud Visualization (DCV) Web Edition (GPU support on HPE ProLiant XL190r Gen9 Server)

Clusters can be customized to add specialty options. The HPE ProLiant XL190r has a 2U depth (1/2 width, so two nodes per chassis), which enables support for accelerator and visualization cards. Accelerators drive application performance by offloading compute-intensive portions of the application to the GPU, while the remainder of the job runs on the processors. Hewlett Packard Enterprise and its partners can provide application-specific information on the potential performance impact.

For sites with very specific application use, Hewlett Packard Enterprise, with its software partners, has developed RAs for ANSYS applications. These RAs can assist customers and system architects determine the best design for running those applications.

Engineering learning curve for using HPC

Tapping into HPC for simulation changes the job submittal process. The user may move from an interactive mode for running simulations on the workstation to submitting to the cluster via a batch scheduler. This can be a simple change, as software developers have enhanced user interfaces to simplify running HPC jobs. Middleware suppliers offer user-friendly job portals with support for multiple applications. Application suppliers provide frameworks that support job flow from pre-processing through post-processing. Workshops and seminars provide an opportunity for hands-on training, as do industry and vendor conferences. Online resources also provide information and training for new users.

New users of HPC can benefit from on-site training and support. As noted earlier, Hewlett Packard Enterprise and its partners offer starter services, including training.

Remote visualization

Simulations can produce sizeable output files, which traditionally need to be transferred to a workstation for post-processing and interpretation. This can be time-consuming and a big inhibitor to leverage the scale-out benefits when engineers are remote from the main HPC clusters.

Remote visualization can fully address these issues, by running post-processing applications right in the cluster where the data is produced, removing the need to transfer data. Remote visualization also uses specialized graphic accelerators and high-memory nodes, therefore, providing enough resources for loading even the largest models and process them at top performance. Moreover, remote visualization can reduce the need for expensive, high-end workstations, as processing is shifted to the shared cluster.

Hewlett Packard Enterprise has worked with technology provider, NICE software, to offer qualified remote visualization on HPE Apollo servers. On these nodes, NICE DCV remote desktop technology enables technical computing users to access their interactive applications over a standard network. Engineers and scientists are immediately empowered by leveraging high-end graphics cards, fast I/O performance, and large memory nodes hosted in a public or private cloud. NICE conveniently works on both Linux and Windows, supports a wide variety of NVIDIA® cards, and many leading engineering applications are certified for running remotely with DCV.

Resources

Solutions for your CAE environment have been designed by Hewlett Packard Enterprise and its partners to enable successful HPC deployments with the aim to improve your productivity. Recommended configurations are developed and tested for optimal performance and operational efficiency. Factory-integrated solutions are available from Hewlett Packard Enterprise and its resellers for rapid deployment.

Hewlett Packard Enterprise has dedicated teams of experienced HPC professionals who work with customers and partners to design and deploy optimal solutions for HPC. For example, the Hewlett Packard Enterprise and Intel Center of Excellence is a high-performance computing engineering and solution development center dedicated to providing the best HPC solutions for today and the future. Located in Grenoble (France), the Center of Excellence is a network of pre-sales, benchmarking, marketing, and development specialists dedicated to assisting customers. It provides a toolbox of expertise and services to help ensure that customer's applications are compatible with today's technologies and those in development. In addition, ANSYS, Intel and HPE are sponsoring a "Free Benchmark Test" program for ANSYS customers who are currently running their ANSYS software on workstations and want to see the benefit of moving to a small cluster. For more information: ansys.com/free-hpc-benchmark

Learn more at

hpe.com/info/ansys

ansys.com/hpc

a00005355ENW, March 2018, Rev. 1