Hewlett Packard Enterprise

Technical white paper

Check if the document is available in the language of your choice.

LEVERAGING HIGH-PERFORMANCE COMPUTING FOR CAE WORKLOADS

Move past the workstation to turbocharge engineering productivity

CONTENTS

CHANGING THE LANDSCAPE OF COMPUTER-AIDED ENGINEERING

Simulation-driven prototyping fuels product innovation and helps users explore their real-world behaviors. Recognizing its benefits, the manufacturing industry has increasingly adopted computer-aided engineering (CAE), and more generally virtual prototyping, in product design.

CAE helps enterprises:

- Improve product quality
- Cope with increasingly demanding regulations
- Increase engineering productivity
- Reduce costs and product development times

While the industry acknowledges the importance of CAE, many organizations have not adapted their technology to capitalize on it. In manufacturing, time is business critical. And productivity improvements boost results. However, for many businesses, engineers accept slow performance of key engineering applications thanks to out-of-date workstations.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.

The answer to these productivity-crippling limitations is a shared, multisystem high-performance computing (HPC) resource. It can provide much faster response times for simulations, support larger model sizes, and support engineering collaboration across the organization. HPC users report faster time-to-market, improved productivity, and deeper insight into product designs. Tapping HPC fosters success in beating the competition, meeting the demands of increased regulation, and tightening staffing levels.

Small companies have been slower than larger firms to adopt HPC and advanced modeling and simulation, 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 up-front capital investment cost and a learning curve to overcome for those new to HPC, the cost for an entry HPC cluster is lower than you 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 x86 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 became more affordable and pervasive, tools and applications were commercialized—with simpler user interfaces and access to professional support services.

We saw this change especially in CAE. Aerospace and nuclear science organizations 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 simulation for the 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, density optimized HPC system can deliver powerful performance in a very small form factor while also easily integrating with existing enterprise servers. For example, a four-node HPE Apollo 2000 Gen10 System fits into standard IT racks, takes 2U of space, and provides the cooling and power efficiencies found with shared infrastructure. With the most current multicore processors, such a system could have 128 2.8 GHz processing cores, 2.0 TB of memory per tray for a maximum of 8.0 TB, and minimal demands on your power and cooling infrastructure.

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.

An HPC system can also be easily (and remotely) accessed by IT or engineering for management, performance, and maintenance. 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 HPC. 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. 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. Because complex models place demands on storage and memory, product designers may constrain the size and quantity of simulations in order to reduce wait times. You can often improve cost and management of multiple high-end workstations 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.

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, reports indicate that analysts are always eager to tap into new technology and compute capacity to run larger models.

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.

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 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 Gen10 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 the number of physical test models and operating costs for the test facilities. In aerospace, wind tunnel 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.

BENEFITS OF MOVING FROM WORKSTATIONS TO HPC

The Hewlett Packard Enterprise computer-aided engineering (CAE) solutions—anchored in HPE Apollo high-performance servers and storage—brings scalability to the world of CAE and analytics packages. HPE's CAE offering maximizes value for manufacturers. It integrates processors, system software, networks, storage, and visualization tools. It leads to quicker deployments and faster time to value (see Table 1).

TABLE 1. Solution platforms: features and benefits

Feature	Benefits
High performance	Reduces time to results, can solve larger and more complex problems with greater accuracy
Efficiency	Provides more performance in a smaller footprint, reduces data center floor space requirements, and lowers energy costs
Reliability	Minimizes downtime and unplanned outages, and improves system availability for engineers and time to market
Scalability	Update or upgrade systems incrementally without replacing large quantities of hardware, dynamically adapts to changing business requirements, streamlines budget planning, and improves predictability of expenses
Remote Visualization	Enhances security by keeping critical data within the data center; boosts productivity and collaboration with anytime, anywhere access to graphic-intensive models; lowers costs by centralization, improving system manageability and optimizing resource (GPUs, software licenses, hardware, and so on) utilization; and promotes retention of highly skilled staff with better work-life balance and location flexibility

Initial investment cost

Entry into HPC can be supported with systems as small as two to four nodes. 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. Sites anticipating expansion or desire to deploy systems as a free-standing resource may use a dedicated rack. The HPE Apollo 2000 Gen10 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 Gen10 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 Performance Cluster Manager—a fully integrated system management solution offering customers all the functionalities they need to manage your Linux-based systems. The software provides system setup, hardware monitoring and management, health management, image management and software updates as well as power management for systems of any scale. The HPE Performance Cluster Manager reduces the time and resources spent administering the system—lowering total cost of ownership, increasing productivity, and providing a better return on your hardware investments. 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 startup 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.

CAE applications

As a trusted advisor, HPE has excellent relationships with independent software vendors (ISVs). Major CAE applications supported and optimized include Altair RADIOSS and Altair FEKO; Ansys Fluent, Ansys Mechanical, Ansys LS-DYNA and Ansys HFSS; MSC Nastran; Siemens STAR-CCM+; OpenFOAM™ and SIMULIA® Abaqus FEA®; and ESI PAM-CRASH to name just a few.

In addition, HPE has computer scientists who help ISVs test and optimize their applications on HPE platforms. Clients benefit from the earliest releases of new versions of application software, as well as optimized versions developed by HPE specialists.

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 2. 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 Gen10 system 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 32 cores per compute node, Intel® Xeon® Gold 6242 2.8 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 single rotor, and 2 x 1600W power supply
Operating system	64-bit Linux and Windows
Cluster management	HPE Performance Cluster Manager (HPECM)
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 Gen10 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.

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 process 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 shifts 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 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 offer a HPC Cluster Appliance Program and 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/hpc-manufacturing-and-engineering

a00005355ENW, March 2020, Rev. 2

Hewlett Packard Enterprise