



**Hewlett Packard
Enterprise**



Sales guide
For HPE and channel partner internal use only

HPC-based Computer- Aided Engineering (CAE): A solutions guide



Table of contents

- 2 **Executive Summary**
- 3 **Introducing Computer-Aided Engineering (CAE) on HPE's high-performance computing (HPC) servers**
- 3 Why is CAE so important to customers?
- 3 Why HPE for CAE?
- 4 What are the leading CAE applications?
- 4 What is the general roadmap for selling an HPC CAE system?
- 5 Lead Generation - Where can one find potential CAE customers?
- 5 What do CAE customers want?
- 6 How can one best identify the key decision makers and who should you call on?
- 7 **Asking your prospect the qualifying questions**
- 7 Qualifying questions for the engineering buyer
- 9 **Developing detailed configuration and proposal — proving you have the solution**
- 9 Hardware stack, software stack, interconnect: The extensive menu of HPE choices for HPC-based CAE
- 13 Assembling your proposed HPC-based CAE configuration
- 14 **Conclusion**

Executive Summary

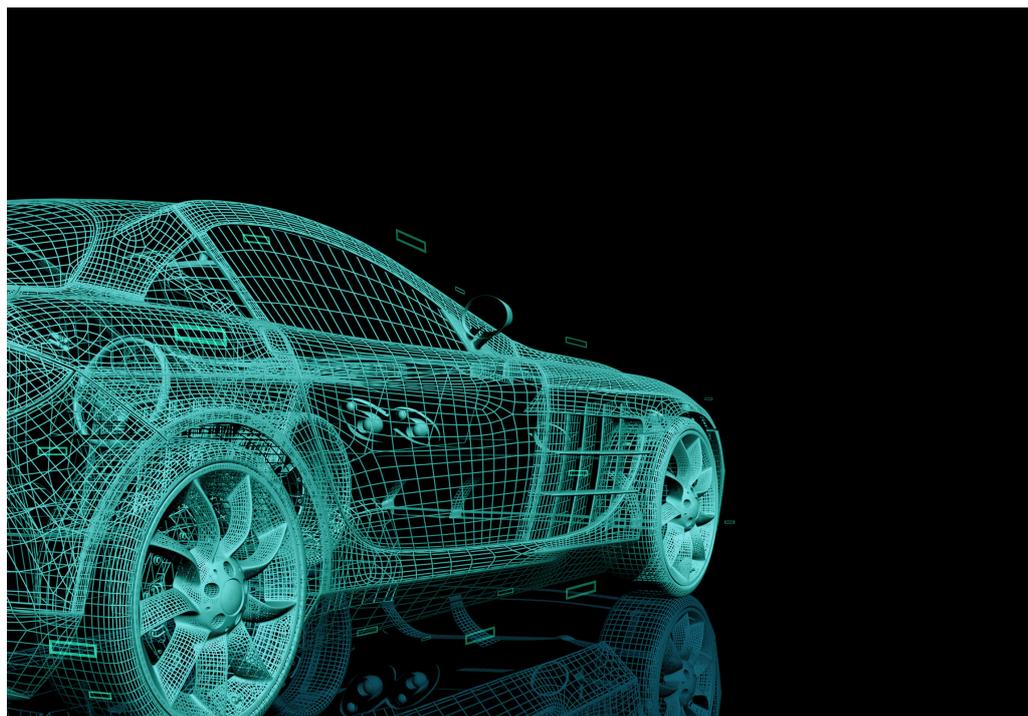
For manufacturers in nearly any competitive marketplace today, the economic benefit of Computer-Aided Engineering (CAE) is enormous. In most settings with an appropriate and customized high-performance computing (HPC) deployment, CAE pays for itself, sometimes several times over. Its ROI, in other words, can be 100% or more.

The Hewlett Packard Enterprise CAE solutions portfolio – anchored in HPE Apollo high-performance servers with Intel® component-level technologies – brings scalability to the world of CAE and analytics packages. HPE’s CAE offering, simply put, maximizes value for manufacturers. It integrates processors, system software, networks and storage. It also leads to quicker deployments and faster time to value.

This document is a guide for our HPC sales organizations and HPC channel partners. It discusses where to find CAE sales opportunities. It also considers how to qualify, position, configure, and sell HPE HPC solutions for CAE. It will help you find and close more—and larger—CAE opportunities.

Remember: Engineering-driven clients most respect and respond not to your “selling” to them; rather, they value your serving as a trusted advisor and partner who understands their pain points and can help solve them. A firm grasp of the material in this paper will give you the grounding necessary to understand the customer’s issues and what unique solutions Hewlett Packard Enterprise has to offer.

This document will be divided into three sections. First, we’ll cover some basic questions about CAE and HPC. Next, we’ll outline the qualifying questions you should cover in your conversations with your potential customer. Finally, we’ll discuss the detailed configuration and proposal you’ll be developing after obtaining answers to your qualifying questions.





Additional Resources:

- [High Performance Computing for Manufacturing and Engineering](#)
- [Engineering with Excellence: CAE Solutions for Manufacturing](#)
- [High Performance Computing and CAE Addressing Challenges in Product Design](#)
- [Turbocharging CAE/CAD Workloads with HPC Modeling and Simulation](#)

1. Introducing Computer-Aided Engineering (CAE) on HPE's high-performance computing (HPC) servers

Why is CAE so important to customers?

Fundamentally, CAE allows customers to design and test ideas for new products without having to build physical prototypes. Virtual prototypes, of course, are far cheaper and quicker to make than physical ones. So CAE typically translates to substantial cost reductions and workflow speedups. CAE also allows companies to focus on only those designs that have the best potential for success. It reduces corporate risk by finding issues early in the design cycle before they go to manufacturing. Other “downstream” effects include reduction of warranty costs and potential litigation if a product failure causes an injury.

A strategic investment in a new, additional or upgraded HPC CAE infrastructure can save the customer hundreds of thousands to millions of dollars in design, manufacturing, warranty, and corporate risk costs. It also can provide additional effects of creating more competitive products, happier customers, and increased revenue for the customer.

Why HPE for CAE?

Hewlett Packard Enterprise provides a comprehensive portfolio of high-performance systems and software as well as high-value services and the best ecosystem of CAE partners to help manufacturing customers reduce costs, improve quality, productivity, and time to market.

Worldwide, many manufacturing companies are already using CAE solutions on Hewlett Packard Enterprise HPC servers.

Hewlett Packard Enterprise delivers a unified CAE compute and storage solution. It's designed to simplify system and data management and to reduce costs and complexity. HPE's market-leading CAE solution also scales to deliver the performance needed for the next-generation of HPC solutions in manufacturing.

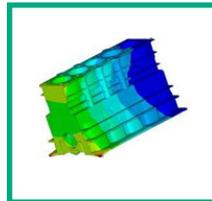




What are the leading CAE applications?

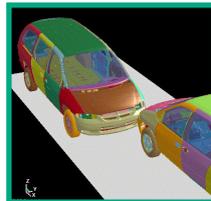
HPC applications generally attempt to model real-world processes or physical systems. In the real world, the process or physical system being observed is often too slow, too fast, too big, too small, or too expensive to observe or prototype empirically. So, with virtual prototyping using CAE applications on HPC systems opens up these heretofore inaccessible worlds to the engineer.

CAE solutions cover a wide range of scientific disciplines, but broadly they span these four application categories:



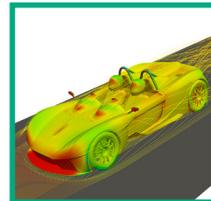
Computational Structural Mechanics (CSM) For Implicit FEA

Simulate the strength and vibration characteristics of a product.



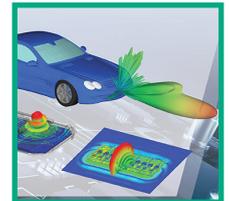
Computational Structural Mechanics (CSM) For Explicit FEA

Simulate the shock impact of products over a short duration.



Computational Fluid Dynamics (CFD)

Simulate Aerodynamics; cooling; mixing of fluids such as air, water and chemicals.

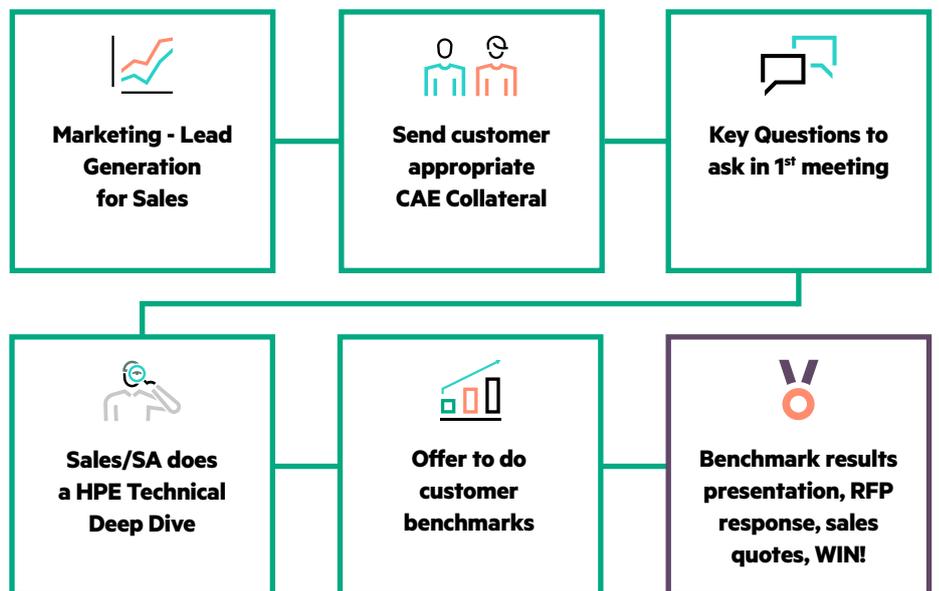


Computational Electromagnetics (CEM)

Simulate radar signature/scattering to assess/prevent detection and identification; Antenna performance, ASIC package simulations.

What is the general roadmap for selling an HPC CAE system?

There are generally six steps in the process toward a successful sale and a satisfied customer. See the flowchart below.



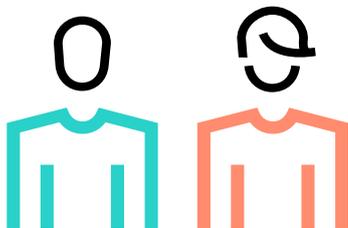
Sales tip

These organizations are always scrambling for cost-effective solutions in order to remain competitive. HPC hardware acquisition can be anywhere during or in-between product design cycles—obsolete equipment represents a performance liability relative to the competition. There’s always an opportunity to help the client if there’s budget available.

CAE solutions usually involve simulating large systems of mathematically defined models, and generally the larger the model, the “better” the simulation depicts the real-world product performance. Compute resources required for the mathematical computations go up exponentially with the size of the model (i.e., if the model doubles in size, it takes four times the compute resources to process it).

CAE customers need:

1. Lower manufacturing and product costs
2. Shortened design cycle times
3. Increase product performance and quality



Lead Generation-Where can one find potential CAE customers?

Organizations that employ HPC for CAE include:

- Automotive Companies and their Tier 1 suppliers
- Aerospace Companies (both commercial and military) and their Tier 1 suppliers
- Consumer Goods/Medical Devices
- Materials and Chemicals Processing
- Electronics and Semiconductor
- Heavy Equipment and Rotating Machinery
- Train and Boat Manufacturers

What do CAE customers want?

While the prospect will undoubtedly have a clear image of their pain points (time to market, product quality, etc.), there are additional benefits to an HPC solution that can help gain their interest early in the sales cycle. These include:

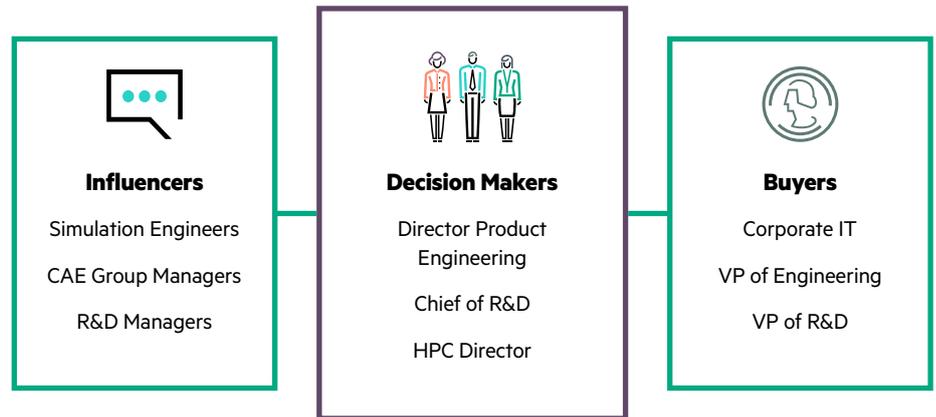
- The earlier, in the product lifecycle, that problems are identified, the less it costs to fix them.
- More simulations and more accurate models result in more competitive and better products. In other words, a productive HPC CAE workflow leads to products that last longer, are less expensive to build, easier to manufacture, and perform better.
- Investments in HPC resources can easily pay for themselves through lower warranty costs, increased product sales, increased regulatory compliancy, higher customer satisfaction, and reduced risk of litigation.
- HPC systems are typically bought with a much shorter financial depreciation model, such as three years instead of five years, meaning hardware upgrades happen more frequently than traditional IT systems.





How can one best identify the key decision makers and who should you call on?

While it's important to include IT during the sales process, the end decision usually rests with the engineering users and their management.



Start with decision makers in upper management on the engineering side of the house, as opposed to the IT side.

It's also productive to meet early on with the influencers. After all, their jobs are most on the line to get products developed on time and on budget, as well as to meet the quality, reliability, and regulatory specifications handed to them. In addition, influencers will be most familiar with their existing suite of CAE hardware and software tools — along with all the existing system's strengths and weaknesses.

IT management is sometimes involved in the procurement decision. However, they also sometimes defer the decision to the engineering management. CAE systems — HPC or otherwise — must meet the engineering department's criteria after being installed. IT often does not want the responsibility for success or failure based on the procurement choices they make. They invariably need the full buy-in of engineering management.

Another user heavily vested in HPC technology (and often overlooked) is that of corporate risk groups. Increasingly complex and stringent legal requirements have forced corporations to examine laws related to compliance and durability. Investing a few million dollars in compliance adherence upfront can potentially avert hundreds of millions of dollars in lawsuits later in the product lifecycle.





Qualifying questions topic areas for the engineering buyer

 **Applications Used and Licensing Needs**

 **Current Workflow and Productivity**

 **Pre and Post Processing Requirements**

 **Compute, I/O and Storage Requirements**


Hewlett Packard Enterprise

Expertise



2. Asking your prospect the qualifying questions

Now that we know how and why HPC-based CAE is so appealing to so many manufacturers today — and we've identified the key decision makers to approach, it's time to formulate the qualifying questions you're going to discuss with your prospect at the first meeting. The answers you get to these questions will help provide the information you'll need to complete Section III, concerning the detailed configuration and proposal that you'll be sending back to this prospect.

Every opportunity is different, but overleaf are 23 general qualifying questions concerning four key areas of the HPC solution.

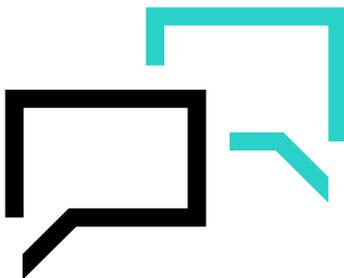
The purpose of asking these qualifying questions is to:

- Understand their current CAE simulation workloads and pain points,
- Understand their CAE simulation goals and where they want to be in 2-3 years,
- Expose and leverage HPE's strategic relationships with ISVs who develop the key commercial applications in use by the customer.

Customer benchmarks

We encourage you to offer a custom benchmark. Each customer has specific models and resolutions (sizes) that can be benchmarked by HPE's CAE Application Engineers, providing a sense of performance gains an HPC CAE solution might provide. Determine the optimal price/performance ratio by running the same benchmarks for several different Intel® Xeon processor SKU's. Benchmarking shows the customer our understanding of the key ISV applications.

To request a customer benchmark, complete an **HPC Sales Opportunity Request form**. Requests for engineering assistance should be submitted with at least two weeks lead time. Very short lead time or emergency requests have a high possibility of being rejected due to resources already being fully committed. Large and/or complex requests will require additional lead time and advance notification.



Applications Used and Licensing Needs

- What are the key simulation applications in use and their order of importance?
- Are there new applications under consideration?
- Are other applications being phased out?
- Are studies periodically done to assess if these applications are being used efficiently and exploiting system capabilities effectively?
- Are there constraints regarding software licensing and licensing costs?

Current Workflow and Productivity

- How many engineers in the company are currently running CAE simulations? Is this number expanding?
- What are their key factors that drive simulation capabilities and requirements? Is it turnaround time for each simulation or number of simulations that can be run each day?
- How is system time and access prioritized between engineers and different departments?
- Are there considerations in job scheduling for interactive job turn around during working hours? What about during non-working hours?
- What are current typical model sizes and simulation lengths for the different solvers in use?
- How many cores do you allocate for various jobs?
- What is the mix between implicit and explicit simulations (where implicit workloads typically require more memory and fast scratch storage)?

Pre and Post Processing Requirements

- Is all your CAD modelling, mesh generation and visualization done on workstations?
- If not, are remote visualization tools used?
- What is the typical size of simulation output your engineers have to do post processing on?
- What applications are used for pre and post processing?
- Are workstations that you do pre/post work on also used to run solvers?

Compute, I/O and Storage Requirements

- Can you please provide the details of your current compute and storage systems?
- Are there any GPU accelerated nodes?
- Are fat nodes (memory rich) or other workload specific nodes deployed/required, i.e. large memory, JBOD local scratch, etc.?
- What is the extent of in house Linux OS, Windows® OS and cluster management expertise?
- Are there parallel file systems in use, i.e. Lustre, GPFS, any others?
- Can you please provide a description of your company's simulation data management and data retention policies?





3. Developing detailed configuration and proposal — proving you have the solution

Hardware stack, software stack, interconnect: The extensive menu of HPE choices for HPC-based CAE

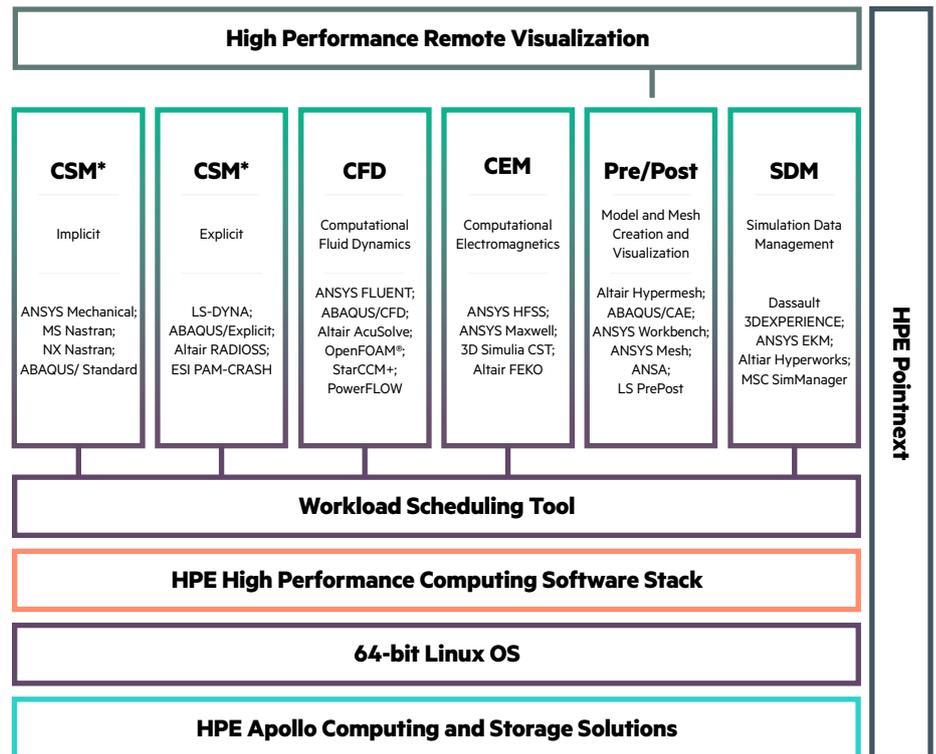
In general, commercial application software drives most deals in the CAE markets. The top 50 commercial applications belong to approximately 10-15 key CAE ISVs. Customers are heavily dependent upon their application software; everything else (platform, operating environment, etc.) is optimized to support the effective and efficient use of these applications.

Hewlett Packard Enterprise has a broad range of platforms, the largest portfolio of application software, and excellent relationships with ISVs. HPE has application engineers who help ISVs develop and tune their codes on HPE platforms, providing our customers with the earliest releases of new versions of applications software, as well as optimized versions developed by HPE specialists.

Solution Environment and CAE Application Segments

HPE Technical Services

- HPE Hardware
- Third Party Software
- HPE Software
- HPE Services



* CSM is Computational Structural Mechanics



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, and ANSYS HFSS; MSC Nastran; Siemens CD-adapco STAR- CCM+; LSTC LS-DYNA; OpenFOAM and SIMULIA ABAQUS FEA; and ESI PAM-CRASH to name just a few.

Addressing CAE Workloads

Because of the range of platforms, HPE salespeople can precisely tailor solutions to fit customer requirements. A tailored HPE solution provides a combination of interoperable platforms that best meets the customers' demands. The components of an HPC solution are discussed in the following sections.

Hardware platforms:

Hewlett Packard Enterprise provides a choice of servers, processors, operating systems, and interconnects—including the HPE Apollo family of servers. The Apollo 2000 Gen 10 is 2U and can hold up to four nodes. The Apollo 6000 Gen10 12U rackmount air-cooled chassis holds up to 24 front-accessible nodes (with up to 3 chassis per rack), or modular HPE SGI 8600, water-cooled compute platform with up to 144 nodes per rack. These clusters combine the flexibility of a custom solution with the simplicity, reliability, and value of a preconfigured, factory-built product. A wide range of qualified options ensures flexible choices, simple implementation, and successful results.




HPE Apollo 2000 Gen10
Mid Range HPC
Flexible scale-out



HPE Apollo 6000 Gen10
Extreme Compute
Performance in
High Density



HPE SGI 8600
Liquid cooled,
delivering industry
leading performance,
density, & efficiency



HPE Superdome Flex
Scale-up, shared
memory HPC, UV
Technologies



Apollo 6500 Gen10
Rack-scale GPU
computing with up
to 8 GPUs per
compute node



HPE Apollo 4510 Gen10
Large-scale, storage
virtualization & tiered
data management
platform

Interconnects:

Hewlett Packard Enterprise supports all major commercial interconnect technologies, including Ethernet, InfiniBand®, and the Intel® Omni-Path Architecture, giving customers a flexible choice to optimize price and performance.

Intel's OPA delivers 100 Gbps port bandwidth with low latency, even at extreme scale. Its 48-port radix switch silicon can reduce the number of switches by as much as 50 per cent in a typical fat tree configuration. This lighter interconnect footprint means less rack space (up to 79 per cent less compared to the competition) and lower power and cooling costs (as much as 60 per cent lower). It also means a reduced fabric cost: a typical OPA solution represents just 21 per cent of an HPC hardware budget compared to the 40 per cent budgetary draw some HPC fabrics represent.

InfiniBand (or IB) is an industry standard for a high-bandwidth, low-latency fabric. Ethernet can be used when you are connecting only two server nodes. For CAE jobs needed to scale greater than two nodes using InfiniBand or Omni-Path interconnect is recommended.



HPE High Performance Computing Software Stack



The HPE cluster software stack that supports modelling and simulation applications typically includes the following components:

The server operating system (including configuration-specific added drivers)

- Cluster tools for management and administration
- HPC libraries and tools
- Job schedulers and resource managers
- Applications

Hewlett Packard Enterprise offers a broad choice for customers across these components, from open source, Linux-based software to proprietary, commercial options. These options include:

Operating systems

In addition to the standard Windows and Linux server options available across ProLiant systems, Hewlett Packard Enterprise offers specially priced and packaged HPC subscriptions for **Red Hat® Enterprise Linux Server** and **SUSE Linux**.

Cluster management tools

Cluster management tools are necessary to install or provision the operating system across the cluster nodes. Other important functions include monitoring overall and server-specific performance, and system administration for server or groups of servers.

HPE Apollo Platform Manager (HPE APM)

HPE Apollo Platform Manager provides the visibility necessary to identify stranded power capacity, ensure safe and effective use of shared server infrastructure, and significantly reduce total data center energy consumption.

HPE Performance Cluster Manager

HPE Performance Cluster Manager delivers an integrated system management solution for Linux®-based high performance computing (HPC) clusters. It provides complete provisioning, management, and monitoring for clusters scaling to 100,000 nodes. The software enables fast system setup from bare-metal, comprehensive hardware monitoring and management, image management, software updates and power management.

HPE Insight Online

HPE Insight Online enables monitoring IT devices from anywhere, at any time, so administrators can work smarter and stay in control of their IT infrastructure.

HPC libraries and tools

Customers typically will acquire and deploy different libraries and tools based on application needs and their preferences — including HPE-MPI, Intel MPI, Intel Trace Analyzer and Collector, TotalView, Allinea DDT and Performance Report, TAU Performance System, and the ScaleMP vSMP virtualization software.



For HPE and channel partner internal use only

For more information on the list of Advisory and/or Professional services listed below, please use the **HPE Pointnext Advisory & Professional Services Knowledge Directory located [here](#)** and search on the name of the advisory or professional service listed below.

Advisory

- HPC Solution Design Workshop²
- HPC Detailed Design Review²
- HPC Redesign Service²
- HPC Technology Adoption Service²
- HPC Custom Consulting Service²
- HPC Backup and Recovery Design²
- HPC for Engineering²
- HPC for Life Sciences²
- HPC for Earth Science²
- HPC Transformation Service²

Professional

- HPC System Tuning Service²
- HPC Job Scheduling Service¹
- HPC File System Services¹
- HPC Resident Service²
- HPC Visualization Services²
- HPC Backup and Recovery Design²
- HPC CXFS Services¹
- HPC DMF Services¹
- HPC Application Tuning Service²
- HPC Dependency Mapping Service²
- HPC Program Management Service²

Helping customers on their unique HPC journey

¹SKU-based packaged services

²SOW-based services

Job schedulers and resource managers

Job schedulers manage the allocation of submitted jobs to nodes across the cluster. These include HPE Performance Cluster Manager, to install specific images on nodes to support different applications, and Altair PBS Professional and Adaptive Computing Moab.

HPE Performance Cluster Manager Connector applications include:

- The Mellanox Unified Fabric Manager (This is a network manager for InfiniBand.)
- The Univa Grid Engine scheduler

Remote Visualization

Remote Visualization: HP Remote Graphics Software (RGS), with high-performance compression, allows users to securely access graphics-rich applications located on a remote system and collaborate in real time from any PC, thin client, or tablet.

Hewlett Packard Enterprise also offers NICE Desktop Cloud Visualization (DCV) through a web portal. DCV provides users with efficient and optimized remote access to graphic-intensive 3D applications including all the major CAE pre and post processing software.

HPC Services (Advisory and Professional)

With the recent acquisition of SGI, 100 additional HPC Advisory and Professionals Services personal have been added to the HPE Pointnext organization. Averaging 18 years' experience of HPC focused solutions, their technical expertise include:

- Servers
- Storage
- Application tuning and optimization
- Visualization
- Networks

HPE HPC services help you to accelerate the HPC solution adoption, empower the platform and client, and reduce the operation burden of the cluster.

Hybrid HPC

HPE is committed to delivering a better compute experience with the right mix of on- and off-premises solutions—including powerful HPC hardware, software, and services. To that end, Hybrid HPC is a critical initiative for HPE. Hybrid HPC delivers enhanced simplicity, security, governance, and control to manage evolving business requirements.

Migrating to a hybrid infrastructure allows organizations to tap into existing on-premises resources as well as scale resources on-demand in the cloud. Hybrid environments are transforming business operations with the right mix of traditional IT and HPC cloud technologies to handle any workload. These solutions create a foundation for a flexible and automated Infrastructure as a Service (IaaS) offering for Hybrid HPC.



Hybrid HPC choice of building blocks:

Options

 Standard HPC cluster with Apollo technology and physical tenant isolation

+

 Hybrid delivery on premises and off premises with and through HPE Partner Ready Service Providers and System Integrators

+

 Choice of HPC automated lifecycle management (built-to-purpose with HPE Performance Cluster Manager and OpenStack®) including Linux containers deployment and PaaS/SaaS connectivity*

+

 Flexible Capacity model for pay per use

+

 Designed, integrated and supported by HPE Pointnext Services

Strategic Program Office

If the CAE customer opportunity is above \$5M please bring this prospect to the attention of the Strategic Program Office (SPO) where you will be assigned a geo-based Capture Manager to help you through the process. The SPO is a new function within the HPC Business Unit.

Assembling your proposed HPC-based CAE configuration

System design and specifications

CAE 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 CAE applications, a well-balanced 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 for CAE Workflows

A typical starter cluster configuration allowing expansion up to four nodes within a 2U chassis 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 ProLiant DL360 Gen10 head node (external) or a single XL170r (within the Apollo 2000 chassis)
Compute nodes	Two to four HPE ProLiant XL1x0r nodes One HPE Apollo 2000 holds up to four XL170r nodes or two XL190r nodes. Customers may also mix and match based on their application needs. Up to 40 cores per compute node, Xeon® Gold 6148 20 core 2.4 GHz processors recommended.
Administration network (out-of-band and console)	One HPE Networking 2920 with 10GbE uplink
Interconnect	10 Gigabit Ethernet, InfiniBand or Omni-Path Architecture
HPE Apollo r2000 chassis	One Apollo r2x00 chassis, each with 4 fans, and 2x1400 W power supply
Operating system	64-bit Linux and Windows
Cluster management	HPE Performance Cluster Manager
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.



Resources:**External:**

[HPE market-leading HPC Solutions for Manufacturing](#)

[HPE ANSYS Solution White Paper](#)

[High Value Computer-Aided Engineering Solutions from Hewlett Packard Enterprise \(HPE\) – a Quantitative Assessment](#)

- [Video Presentation](#)

[Engineer with excellence LS-DYNA Solution for CAE](#)

[Engineer with excellence HPE ANSYS Solution for CAE](#)

[Total CAE ANSYS Apollo 2000 Gen10 Starter-Kit](#)

[Engineer with excellence Simulia Solution for CAE](#)

Internal:

[HPE ANSYS quick reference card](#)

[HPE ANSYS competitive Battle card](#)

[HPE Customer Benchmark Request Form](#)

Conclusion

The computer-aided engineering (CAE) segment of HPC represents a huge potential for sales. Not only are the prospects in need of HPE's products and services to remain competitive, but they are also incentivized to remain with a vendor once chosen. Gaining a foothold in a large and medium manufacturing account can be immediately lucrative, but also represents a potential stream of sales for years to come.

Our commitment, products, and services to customers in these industries represent a highly probable win—creating satisfied customers who regularly return to refresh their technology and can help sell to additional accounts.

People and Organizations

Tony DeVarco	HPE, Manufacturing Vertical and GTM Manager HPC/AI BU anthony.devarco@hpe.com
Liz King	VP HPC Sales Enablement and GTM HPC/AI BU Acting Head of the HPC Strategic Program Office (SPO) elizabeth.king@hpe.com
Brian Thatch	Director, HPC Applications Engineering and Benchmarking brian.thatch@hpe.com
Jean-Luc Assor	Hybrid HPC Program Worldwide Lead HPC/AI BU jean-luc.assor@hpe.com
Ramesh Joginpalli	Hybrid HPC Program Manager HPC/AI BU ramesh.joginpalli@hpe.com
Everton Paulino	AMS HPC Sales Consultant Everton.j.paulino@hpe.com
Keith Mullin	AMS HPC Sales Consultant Keith.mullin@hpe.com
Kevin Denton	Manager, HPC Services Pointnext Kevin.denton@hpe.com
Ghislain De Jacquelot	EMEA HPC North and West Sales Ghislain.de-jacquelot@hpe.com
Jose Rodrigues	EMEA France, East and South Sales jose.rodrigues@hpe.com
Michal Klimes	EMEA HPC Channel Partners Michal.klimes@hpe.com
Florian Jay	EMEA HPC Marketing Florian.jay@hpe.com
Dan Boone	Senior Manager, EMEA HPC Services Pointnext Daniel.boone@hpe.com
Grenoble Competency Center (including all COE's)	hpc-emea@hpe.com
Hiroaki Kondo	APJ HPC Sales Director Hiroaki.kondo@hpe.com
Laurent Herviou	APAC HPC Sales Director Laurent.herviou@hpe.com

 Share with colleagues

© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

This document contains confidential and/or legally privileged information. It is intended for Hewlett Packard Enterprise and Channel Partner Internal Use only. If you are not an intended recipient as identified on the front cover of this document, you are strictly prohibited from reviewing, redistributing, disseminating, or in any other way using or relying on the contents of this document. Trademark acknowledgments if needed. All other third-party trademark(s) is/are the property of their respective owner(s).

Intel, Intel Xeon, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries. Windows is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. All other third-party trademark(s) is/are property of their respective owner(s).
a00048487enw, June 2018