


Deep Computing Capacity on Demand: Automotive Industry



Highlights

- *Addresses peak workload demands for high performance computing*
- *Offers flexible pricing, allowing customers to pay only for the capacity reserved and helps avoid up-front capital investment*
- *Helps create a flexible, scalable and responsive IT infrastructure*

IBM Deep Computing

IBM Deep Computing can address the needs of customers who require powerful solutions to address their most complex and challenging business needs. These customers include scientists, engineers, knowledge workers, and other professionals whose core business or research relies on high performance computing (HPC) for deep levels of computation, data access, visualization, and communications. Deep Computing is more than traditional HPC hardware, software and services. It also embraces emerging strategic technologies, research initiatives, open standards, industry and customer partnerships,

and industry expertise. IBM Deep Computing takes a comprehensive approach to helping our customers gain insight and transform their business and organizations. Industries and business segments include Life Sciences, Digital Media, Financial Services, Petroleum, Electronics, Automotive, Aerospace, Government and Higher Education.

Using HPC in the Automotive industry

In a race where dramatic reductions in concept to launch time determine who wins, the effective use of analysis, simulation and optimization technology can make a difference.

Automotive vehicle manufacturers and their suppliers are taking advantage of analysis, simulation, and optimization solutions in the computer-aided engineering (CAE) areas of crash and stamping simulation, computational fluid dynamics including aerodynamics, and structural analysis including noise, vibration and harshness, to get their products to market with less physical prototyping.

High Performance Computing

Using less physical prototypes allows auto makers to get their products to market faster and at less cost. To support the increased reliance on automotive analysis and simulation solutions, a high performance computing (HPC) infrastructure is required to deliver fast feedback for design optimization early in the design phase. However, if the HPC environment has not been designed to meet the demands of overlapping vehicle development programs, project launch dates may slip and the costs associated with last minute changes may increase.

IBM's offerings for on demand delivery of high performance computing can help vehicle manufacturers and suppliers meet demands for peak capacity, while providing an alternative approach to managing capital and operating costs. Automotive companies can achieve attractive price/performance for compute-intensive applications and processing of numerous large data files. These companies may also reduce the administrative costs associated with managing distributed systems and data, allowing them to focus on their core business functions rather than on IT cost recovery or infrastructure complexity.

Meeting automotive business challenges

If one could select a single adjective to describe the automotive industry, the most apt descriptor would be: complex. Market segments are large and diverse, from passenger and light vehicles to heavy trucks and equipment. Industry value nets are comprised of intricate webs of suppliers, auto makers and retail dealerships. Not to mention that automotive is a global game, with design, manufacturing and distribution occurring on virtually every continent.

The intense competition for market share has generated a frenzy of mergers and acquisitions that left frequently larger companies complex, if not unwieldy, organizations and infrastructures. To top things off, the innovation race has multiplied the technological complexity of vehicles and the processes used to design, validate, produce and service them. This complexity increases the possibility of longer time to market, higher lifecycle costs and more quality problems.

To keep business opportunities from getting mired in all of this complexity, automotive companies need a different approach; one that allows them to cost effectively produce the innovation demanded by customers, team with partners to drive down development time and costs, and enter new markets when and where appropriate.

Automotive companies are using more computer-aided engineering methods to engineer and refine their new vehicles without the use of costly prototypes. A flexible and powerful high performance computing infrastructure enables automotive companies to design and deliver high quality products, while managing development costs, vehicle costs and time to market.

Automotive companies strive to meet government regulations and consumer demands for safety features in cars. They accomplish this by using CAE to predict failure points and simulate the vehicle's collision behaviors. This upfront analysis allows the automotive companies to deliver safer vehicles to market sooner by eliminating timely and expensive physical

testing early in the design cycle. The applications which predict crash characteristics are also used by automotive companies to predict the formability of sheet metal for stamping operations in manufacturing. This formability analysis allows companies to shorten their tool development and tryout times. These crash and stamping applications represent the biggest workload for HPC infrastructures.

Computational fluid dynamics is used to predict vehicle characteristics such as wind resistance, wind noise, comfortable temperature zones for passengers, and engine cooling. This type of analysis is faster and less costly in a CAE environment than using physical prototypes.

Finally, the area of structural analysis allows automotive companies to deliver better vehicle ergonomics and comfort to create a more enjoyable driving experience. This is accomplished by analyzing the vehicle design to eliminate squeaks and rattles before they become issues.

Performing this analysis, simulation and optimization, even during periods of high peak usage due to unpredictable demands, can be critical to the continued success of automotive companies.

The automotive industry is facing voluntary and mandatory process and model transformations that are driving an increasing, unpredictable and variable demand for HPC capacity. Planning against such fluctuating demand becomes inherently difficult. Companies must provide services quickly, efficiently and cost-effectively. The successful companies are those that meet the fluctuating needs of the automotive industry while helping to control costs, advance product development, and improve profit margins. Using IBM cluster services to obtain HPC capacity on demand can contribute to this success.

IBM can help the automotive industry meet these analysis, simulation, and optimization business challenges by

helping to maximize capacity flexibility while minimizing capital investment. IBM offerings for delivering HPC on demand are based on state-of-the-art cluster technologies and services, and leverage IBM's industry and technology expertise.

Examining HPC requirements for the automotive industry

Providing HPC resources on demand requires an infrastructure of high-speed processors and storage that can provide measurable compute and data storage capacity in a flexible, scalable and highly secure manner. Users require seamless access to these resources, regardless of where the physical systems may reside. An HPC infrastructure for the automotive industry has the following objectives:

Scalability: Applications and databases often exceed the capacity of even the largest single system, so systems need to grow with demand.

High Performance Computing

Security: Automotive engineers need a security-rich processing and information environment.

Flexibility: While they focus on controlling costs, automotive companies need increased flexibility in a rapidly changing environment.

Significant processing capacity: As processing requirements continue to grow, more and more companies are finding that they can benefit from using clustered servers and Grid computing. Clusters link separate nodes to create a single, powerful, scalable system, while Grid computing increases overall compute resource sharing and utilization.

Flexible financial and delivery models: Automotive companies can benefit from balancing their fixed and variable IT infrastructure and operations costs and choosing between onsite owned/operated versus offsite hosted IT solutions.

Clustering and Grid computing, combined with a wide range of server and storage technologies, are designed to provide:

- *High availability with failover protection*
- *Scalability with minimal downtime*
- *Large-capacity data volumes*
- *Capability to handle peaks in workload, permitting flexible responses to changing engineering and business requirements*

Providing high performance computing on demand

Fundamentally, providing HPC capacity on demand is a simple concept. Instead of purchasing computers to meet peak demands, companies can procure sufficient hardware for average demand and then contract remote processing power to help meet peak loads, or offload all processing to a remote facility. In practice this process can be more complex. However, IBM can help shield customers from much of this complexity—one of the primary advantages of a hosted on demand environment.

High performance computing can have many different requirements, and purchasing sufficient hardware can become expensive. Some problems are computationally parallel, and a large number of servers working in parallel can solve the problem quickly and efficiently. Other problems require ultrafast interprocess communication or a single large memory model. A hosted environment can supply access to these resources, and instead of paying for equipment that may be underutilized, customers simply pay for the cycles and compute or storage resources that they reserve for a specified period of time.

IBM provides a scalable, highly secure and extensible environment designed to handle peak workloads. This solution combines the best of IBM technologies and open standards with on demand computing, providing an infrastructure that enables customers to help meet their goals of increased scalability, flexibility, and processing capacity.

Building a world class infrastructure to support the automotive industry

IBM has a comprehensive approach to deliver on demand computing for automotive CAE markets. IBM's advanced clustering capabilities and services can help reduce the time, effort and expense required to analyze, simulate and optimize a broad spectrum of automotive design challenges.

The HPC infrastructure is based on the IBM **@server**™ Cluster 1350, an innovative Linux-based cluster that provides a robust, highly scalable compute facility centrally managed and controlled by IBM. The Cluster 1350 combines the power of IBM **@server** xSeries® Intel Xeon™ processor 32-bit and IBM **@server** AMD Opteron™ processor 32-bit/64-bit compute servers with IBM Cluster Systems Management (CSM) software, IBM TotalStorage® products, and leading third-party networking components to enable powerful,

flexible solutions for high performance computing. IBM also offers clusters of IBM **@server** pSeries® POWER™ 64-bit compute nodes running AIX 5L™ or Linux. IBM plans to offer other models, platforms and technologies over time.

Customers can access virtual dedicated Cluster 1350 cluster capacity based on the xSeries 335 2-way Intel Xeon processor-based server or IBM **@server** 325 2-way AMD Opteron processor-based server. The xSeries 345 or IBM **@server** 325 is used as a management node for centralized control of the system. Each virtual cluster has its own dedicated management node. Customers can also access virtual cluster capacity based on the pSeries 655 4-way server with a pSeries management node. Optional Fibre Channel or SCSI external disk storage capacity can be provided. Customers also have access to an IBM 3590, LTO or DLT serially shared tape server for loading data and performing backups.

Customers have full control of the compute and storage resources within their assigned environment, and each virtual cluster has a dedicated cluster and management network. A robust networking infrastructure is designed to allow customer data and applications to be highly available and secure, and a virtual private network (VPN) service provides remote access to the facility. Customers pay for their assigned compute and storage capacity for the length of time that they use it.

Advanced systems management and Grid-enabling software

IBM Cluster Systems Management (CSM) for Linux software provides robust capabilities for resource monitoring, automated operations, remote hardware control, distributed command execution, configuration file management and parallel network installation. Access to the system is accomplished through a remote shell interface. Customers may use CSM directly to perform remote management tasks, install and use their own management tools, or request comprehensive management and monitoring services from IBM.

High Performance Computing

The IBM General Parallel File System (GPFS) for Linux and AIX®, offered as an option, is a high performance scalable file system—ideal for automotive companies that manage extremely large data files. GPFS can provide an enhanced layer of scalability, availability and performance. Optional dynamic job scheduling and workload management tools are also planned to be available to optimize cluster resources and help increase job throughput.

Customers supply the additional software required to run their workloads; this software may include in-house, open source, or IBM and third party applications and tools. Customers can use Grid software and services, acquired separately from IBM or a third party, to develop new applications or run existing applications in a Grid environment. The open source Globus Toolkit™ is designed to provide the underlying elements for Grid security, communication, information infrastructure, resource management

and portability. The IBM Grid Toolbox for Linux and AIX is an integrated set of tools and software based on the Globus Toolkit that facilitates the building of Grids and development of Grid applications. In addition to running production workloads, customers may run benchmarks or perform other types of analysis.

Services for delivering HPC on demand

IBM offers a variety of services to support IBM-supplied hardware, software, and networking components and infrastructure. Base services include the facility, provisioning and deployment of resources, security, virtual private network, monitoring, maintenance, and help desk support. Customers can benefit from a range of fee-based custom services, such as enhanced VPN bandwidth, software deployment and customization, advanced monitoring and management, storage management and backup, Grid-related services, and more.

Delivering dynamic HPC resources

IBM offers dynamic HPC resources on demand that are designed to match customer requirements for scalability and capacity. The on demand computing model is economical, enabling customers to enjoy flexible pricing and terms and helping them to realize lower operational and infrastructure fixed costs. This computing model also can help reduce infrastructure complexity, IT training and staffing and ongoing maintenance and upgrades. Customers can focus on their core engineering issues and be insulated from rapid information technology shifts.

IBM is extending its leadership in high performance computing to the on demand environment. By combining leading IBM @server technologies, advanced cluster and systems management capabilities, growth and innovation in Grid and autonomic computing, and demonstrated industry expertise, IBM is delivering solutions designed to enable automotive customers to conduct their analysis, simulation and optimization more efficiently and effectively in a highly competitive market.

Figure A. Possible technical components: Deep Computing Capacity on Demand for the Automotive Industry

Component	Features/Benefits
HARDWARE	
IBM @server Cluster 1350	<ul style="list-style-type: none"> • Combines IBM @server servers, IBM TotalStorage storage and leading third-party networking components to create powerful, flexible solutions for HPC and commercial application environments • Uses xSeries 335 (1U) and xSeries 345 Intel Xeon 32-bit dual processor nodes • Uses BladeCenter Intel Xeon 32-bit dual processor nodes • Uses IBM @server 325 (1U) AMD Opteron 32-bit/64-bit dual processor nodes • Runs the Linux operating system standard • Runs the Windows operating system by exception • Uses CSM for central management
AMD Opteron Technology	<ul style="list-style-type: none"> • Runs 32-bit and 64-bit applications simultaneously • Easy migration to 64-bit computing capability with software investment protection • Extreme price/performance design—high performance at affordable pricing • Runs Linux and Windows
IBM pSeries Technology	<ul style="list-style-type: none"> • Features POWER 64-bit processors in an ultra-dense packaged server for high performance computing • Includes the AIX 5L operating system and a rich suite of cluster software
IBM TotalStorage	<ul style="list-style-type: none"> • FASTT Fibre Channel and SCSI disk options • 3590, LTO, and DLT tape server and cartridges options
SOFTWARE	
Operating Systems	<ul style="list-style-type: none"> • Linux and Windows provided by customer • AIX provided by IBM
IBM Cluster Systems Management (CSM)	<ul style="list-style-type: none"> • Provides robust, powerful management from a central point of control • Simplifies administrative tasks and may reduce life-cycle costs • Offers a highly reliable infrastructure and event monitoring • Provides software installation and updates, remote hardware control, distributed command execution, configuration file management, and diagnostics
IBM General Parallel File System (GPFS) (optional)	<ul style="list-style-type: none"> • Provides shared access to files across multiple disk drives on multiple nodes • Provides a common file system abstraction for data shared among multiple nodes • Allows applications to easily access files using standard POSIX (Portable Operating System Interface for UNIX®) file system interfaces • Enables parallel applications to simultaneously access either the same or different files • Provides high availability through automatic recovery from node and disk failures
Job scheduling (optional) (planned availability)	<ul style="list-style-type: none"> • Provides dynamic job scheduling and workload management • Designed to optimize cluster resources and increase job throughput • Works with CSM to facilitate management of cluster resources
SERVICES	
IBM Global Services	<ul style="list-style-type: none"> • Support Deep Computing Capacity on Demand with a full range of services for solution customization • Provide Grid enablement services, including Grid Innovation Workshops, Grid Pilot Implementation Services, and Grid Rollout Implementation Services

For more information

To learn more about IBM Deep Computing Capacity on Demand for the Automotive industry, contact your IBM representative.

AMERICAS

Alan McCarter
e-mail: armccar@us.ibm.com

ASIA PACIFIC

Sinisa Nikolic
e-mail: 9sinisa@my.ibm.com

EUROPE/MIDDLE EAST/AFRICA

Ian Green
e-mail: ian_green@uk.ibm.com



© Copyright IBM Corporation 2003

IBM Systems Group
Route 100
Somers, NY 10589

Produced in the United States
December 2003
All Rights Reserved

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, pSeries, TotalStorage, xSeries

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: AIX 5L, @server, POWER, POWER4, POWER4+

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

AMD Opteron is a trademark of AMD.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Globus Toolkit is a trademark of the University of Chicago.

Other company, product and service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All information in these materials is subject to change without notice. ALL INFORMATION IS PROVIDED ON AN "AS IS" BASIS, WITHOUT ANY WARRANTY OF ANY KIND.