IBM Deep Computing

November, 2005

Table of Contents

Executive Summary

Deep Computing offers solutions to previously intractable business and scientific problems, by exploiting computational strengths in high-end computing, data storage and management, algorithms, modelling, simulation, visualization and graphics. Often these solutions provide new knowledge and understanding at the frontiers of science and business. Deep Computing is a refinement of supercomputing or high performance computing (HPC), recognizing there is more to the problem than raw compute power.

In the past, HPC was dominated by vector architecture machines. With the advent of powerful super-scalar processors and the ability to interconnect them in parallel to execute a single job, vector systems were eclipsed by massively parallel RISC machines, which in turn led to clusters of commodity hardware. While parallel RISC machines still satisfy the most demanding application needs, clusters have also proven to be a cost-effective method for many HPC workloads. It is common today for clusters to be used by large corporations as well as small and medium enterprises, universities and government labs to solve problems in life sciences, petroleum exploration, structural design, high-energy physics, finance and securities, and more.

IBM's experience in HPC spans nearly half a century: from early extensions of IBM mainframes to support floating-point operations, to the introduction of Unix® engineering workstations and servers, to the launch of the RS/6000 SP1®, the first viable parallel RISC machine in the industry and a market leader among HPC systems. In the last decade, IBM has leveraged the robust software stack developed for the SP by making it available on Linux®, and introducing a portfolio of Linux cluster offerings that yield excellent price/performance for scale-out computing.

IBM continues to lead the HPC industry with innovative offerings, such as IBM BladeCenter™ and Blue Gene®, which address critical issues of power and space consumption, scaling, integrated networking, and centralized systems management. IBM's strategy in this space is simple: IBM is dedicated to solving the most challenging and complex problems at lower cost, enabling our clients to innovate and gain competitive advantage with innovative technology.

IBM's portfolio of Deep Computing solutions is extensive. IBM is fully committed to the Power Architecture™, and offers systems based on POWER as well as Intel® and AMD® processors in response to a wide range of marketplace demands. POWER is at the foundation of the HPC product line, and forms the basis of special solutions like Blue Gene.

IBM strongly supports open standards, and the entire server product line is enabled to run Linux. The Deep Computing portfolio includes high performance workstations, storage hardware and software, HPC software and tools, on demand and grid computing, and visualization solutions. IBM also delivers integrated solutions to address compute- and data-intensive challenges in areas such as life sciences, digital media, product lifecycle management, and business intelligence. IBM’s investment in Deep Computing is evidenced by its dominant leadership on the Top500 supercomputing list.

IBM recognizes that one single architecture or system design does not address the range and complexity of problems in the HPC market place. Clients want choice, and IBM offers a rich portfolio to address the diverse needs of this market. Get IBM expertise working on your problem today.

Background

Deep Computing delivers powerful solutions to clients' most challenging and complex problems, enabling businesses and researchers to get results faster, formulate a decision based on these results, and realize a competitive advantage. Deep Computing involves the design of novel algorithms that take advantage of the architecture, to gain insight on a problem of interest to society or business. HPC systems provide extreme computational power that is frequently measured by how many Trillions of Floating-Point Operations (teraflops) can be performed every second. Models of extreme complexity can be analyzed to gain insights never before possible, resulting in better informed decisions.

High performance computing has traditionally been characterized by computational speed. Beginning with fast scalar computers, the market was soon dominated by vector architecture machines, which boasted the fastest computation speeds of the time. Users recognized that getting data to the computational engine was equally important, requiring focus on both the computational capability and on managing the flow of data to minimize time to solution. Problems with huge datasets or vast relational databases in addition to extreme computational needs became tractable, although expense often gated the feasibility of solving these problems quickly and affordably.

Then came the superscalar RISC (Reduced Instruction Set Computing) architectures. RISC-based microprocessors in workstation packaging spawned the creation of parallel UNIX machines, using workstation as nodes and interconnecting them with a switch fabric. The HPC community initially viewed the new parallel machines as a novelty, but over time the algorithmic advances allowed the parallelization of many traditional vector and sequential applications, providing increased performance.

While some of the most ambitious HPC work takes place in government laboratories and universities, much is now performed for commercial applications as well. The range of industries and workloads addressed by Deep Computing includes the following:

- Government and academia - research, classified and defense workloads
- Environmental sciences – weather and climate modeling
- Petroleum exploration and production - seismic processing, reservoir simulation
- Automotive and aerospace computer-aided engineering - crash analysis, structural analysis, computational fluid dynamics, noise/vibration/harshness analysis, design optimization
- Product lifecycle management - computer-aided design and product data management
- Life sciences - drug discovery, genomics, proteomics, quantum chemistry, structure-based design, molecular dynamics
- Electronics - design verification and simulation, auto test pattern generation, design rule checking, mask generation, optical proximity correction
- Financial services - Monte Carlo simulations, portfolio and wealth management, risk management, compliance
- Digital media – animation, special effects, rendering, online games, digital security and surveillance, content creation and management
- Business intelligence - data mining, data warehousing, decision support

IBM Deep Computing Heritage

IBM has a rich history of innovation in technical computing, with arguably the longest track record in the industry. Recent dominance in the Top500 list speaks to the depth and breadth of IBM's passion and commitment in this market.

IBM's experience in technical computing goes back to the late 1950's with 7090/7094 systems running FORTRAN compilers. While the System 360/370 was considered primarily a business computer, it was suitable for both commercial and technical applications, with special hardware support for floating-point calculations. In the late 1980's, IBM introduced the first Vector Facility addition to the System 390. At about the same time, IBM introduced an engineering workstation (RS/6000) based on the IBM RISC design called POWER. While the early RS/6000 workstation had only one CPU, it delivered far better graphics and floating-point performance than the PC's of the day, and was quickly embraced by scientists and engineers for all kind of computing needs.

Building on the success of the RS/6000 workstation, IBM began focusing on multiple CPU's in a SMP (symmetric multiprocessor) server, as well as parallel architectures. Scalability was a critical factor in system design, and an IBM Research switch project addressed the scalability problem. As nodes were added to the parallel computer, paths were added in the switch. There was no bus to share, or complex network to navigate via a variable number of hops. Each path was a direct link with no contention. Sounds simple, but it took many years to develop a solution that would be scalable no matter how large the system might be.

IBM changed the HPC landscape in 1993 with the introduction of the RS/6000 SP1, the first sustainable and viable parallel RISC machine in the industry. The SP1 was a completely integrated solution: fast microprocessors, a scalable switch, an integrated software stack, job scheduling, and systems management. Even though each node was a complete system, the software provided a single point of control and a single systems image across all nodes.

In 1997, IBM fielded the famous *Deep Blue* chess playing computer against world chess champion Gary Kasparov, and won. Deep Blue was a modified RS/6000 SP, with special hardware and software that allowed chess playing at the grand master level. No computer of the time could examine all possible moves from a chessboard in a particular state, and within the time allowed for a tournament play. Algorithmic work was needed to decide, in very fast real time, which moves to evaluate and which to ignore.

The complexity of the algorithms needed for Deep Blue, following the complexity of the algorithmic change from vector to parallel, led to the characterization of *Deep Computing*, where computation, data, and algorithmic complexity (or elegance) together challenged the limits of computational capabilities.

Starting with a Department of Energy project in 1994, similar parallel architectures were being custom built using commodity hardware, called Beowulf clusters. The objective was to achieve supercomputing power from clusters of Common Off The Shelf (COTS) components, often using a special high speed network switch for communication. Today's commodity clusters are primarily built of Intel or AMD processors (x86 architecture) and typically run Linux as an operating system.

Conceptually, Linux clusters have many of the same attributes as the RS/6000 SP – densely packaged compute nodes, centralized management and storage nodes, cluster software, a high performance switch fabric – and like the SP, clusters provide most value when they are an integrated solution. Recognizing that many Deep Computing needs could be met by clusters, IBM applied its extensive experience in development and deployment of parallel UNIX machines to the Linux cluster space.

IBM entered this market with an integrated offering, consisting of densely packaged rack mounted servers using Intel Pentium 4 processors with one or two CPUs per node, along with industry standard switches, storage, and cluster software. Over time the solution evolved, becoming the IBM eServer Cluster 1350 in 2003. Around the same time frame, the IBM BladeCenter was introduced, combining one or two CPUs per blade with storage and adapters, in a package twice as dense, and with communication networks and power built into the blade chassis. IBM ported the SP management software from AIX to Linux, introducing IBM Cluster System Management (CSM) and IBM General Parallel File System (GPFS) into the portfolio. Partnerships were formed with suppliers of other parts of the cluster infrastructure, like Myricom® with the Myrinet® switch and TopSpin® with the InfiniBand® switch.

In 1999, IBM Research began a project to design a new supercomputer to address the challenge of protein folding, a foundational element in drug design in the life sciences field. This was the origin of the Blue Gene system. Using Power Architecture as the core, IBM's architects focused on processor and system density, scalability, power consumption, cooling, systems management, reliability, and familiarity of programming environment, as opposed to absolute clock speed and theoretical price/performance. The design goals were teraflops/watt and teraflops/square foot of floor space, instead of the traditional metric of teraflops/$.

In 2004, IBM announced plans to commercialize the Blue Gene system, enabling its use across a broader range of applications and industries. IBM is continuing its pace of innovation and leadership in Deep Computing, extending the capabilities of its server and cluster systems while introducing new solutions and business models to address the supercomputing challenges of today.

IBM Deep Computing Strategy and Portfolio

IBM's strategy for Deep Computing can be articulated simply – IBM is dedicated to solving the most challenging and complex problems at lower cost with innovative technology, to provide business value. Market dynamics and the needs of clients will continue to drive the pace of technology and solution development, and innovation will enable change - both evolutionary and disruptive change. IBM’s objective for Deep Computing is to be solution driven.

In execution of this strategy, IBM is fully committed to the Power Architecture, which is ideally suited to deliver maximum performance and robustness for the highest performance servers, including large SMP’s and Blue Gene. For mid-range and density markets, IBM delivers clusters based on POWER, Intel and AMD processors. IBM will accelerate the use of its considerable technical resources, including IBM Research, to lead the latest industry trends, including clustering, chip technology, and Linux. Extending the value of its clustering capabilities, IBM offers supercomputing resources on a pay-as-you-go basis through its Deep Computing Capacity on Demand centers, which provide access to massive cluster resources as well as Blue Gene through IBM owned and operated centers. The Deep Computing portfolio also encompasses storage solutions, HPC software, grid computing, and visualization solutions.

POWER

Power Architecture has been and continues to be the flagship of IBM's HPC portfolio. Now in its fifth generation (POWER5), the roadmap continues to extend the performance and capabilities of this architecture. Along with the expected improvements in hardware, like faster clocks, more transistors per chip and closer tolerances, there is more internal parallelism. The number of instructions in flight in a single processor at an instant in time is now in the hundreds. Storage reference patterns are noted by the hardware and pre-fetching is initiated when appropriate. Memory subsystems (cache at multiple levels as well as main store) have continued to improve, and critical components of the memory subsystem are integrated on the chip.

The POWER5 system family extends from small SMP’s up to the most powerful SMP and cluster servers. Designed expressly for high performance computing applications as well as business intelligence, the IBM p5 575™ provides the basic building block for POWER clusters. Each p575 has 8 single core POWER5 CPUs, with 36MB of dedicated cache per CPU. As a component of the Cluster 1600, 64 p575 nodes can be attached to an IBM High Performance Switch™ (HPS), for a total of 512 processors. Customers may choose AIX or Linux for the p575. This system is ideally suited for applications such as data warehousing, scientific research, national security, weather forecasting, and automotive design. Mid-range servers based on the POWER architecture support HPC applications which are optimized for smaller SMP’s, including industrial workloads and Product Lifecycle Management (PLM).

Clusters

Clusters using Intel, AMD, or POWER processors continue to meet the needs of many Deep Computing clients, and offer excellent price/performance for “scale-out” applications such as crash simulation, computational fluid dynamics, biosciences, seismic analysis, reservoir simulation, electronic design, financial analytics, and digital animation and rendering. For most clusters, Linux is the operating system of choice. The increasing availability of fast interconnect technologies make clusters viable for an even larger marketplace.

Cluster 1600

The IBM eServer Cluster 1600, configured with high performing POWER5 servers and the IBM High Performance Switch, delivers extreme performance for the most demanding applications. Nodes of the Cluster 1600 range from the p575 up to the largest SMP's in the pSeries family, the p595 with up to 64 CPUs.

The IBM eServer Cluster 1350 utilizes dense rack packaging (typically 1U or 2U servers with 2 or 4 CPUs). Storage, switches, and infrastructure are included in the rack. The IBM Cluster 1350 is a fully integrated solution, tailored to clients' needs. IBM offers a wrap-around service agreement, covering support of all the cluster hardware and software.

The Cluster 1350 can also include BladeCenter nodes, offering effectively twice the CPU density, with infrastructure like power, cooling, and I/O provided by the chassis. Blades may use Intel, AMD, or POWER processors. The JS20 incorporates two-way Power PC® processors running Linux, for optimal HPC performance.

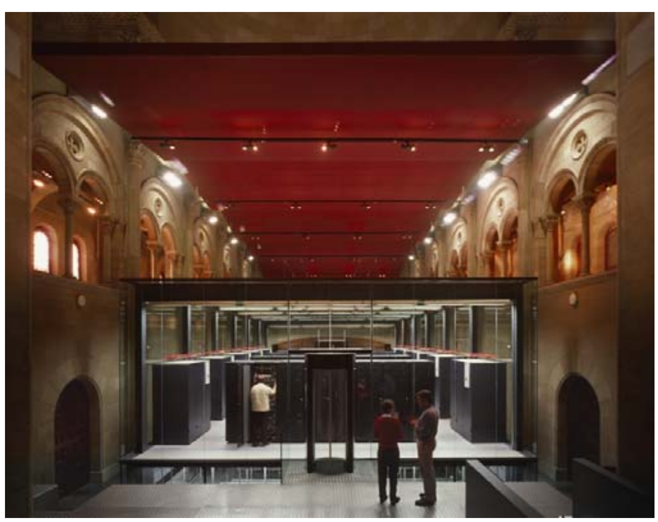

Mare Nostrum at University of Barcelona Supercomputing Centre, Spain
BladeCenter using JS20 Power Blades

On Demand

Many IT workloads have both predictable and unpredictable components. Unscheduled new projects or expansion of existing projects are a reality. The traditional solution of buying enough capacity to cover the peaks, while effective, can become an expensive proposition. One alternative is to install just enough capacity in the client's datacenter to cover the minimum or average workloads. When additional capacity is needed, the client can then access the supercomputing power hosted by IBM, using a highly secure Virtual Private Network (VPN) over the Internet.

This model delivers significant advantages for Deep Computing clients, who typically have fixed capacity at a fixed cost. IBM's Deep Computing Capacity on Demand (DCCoD) centers provide variable capacity at a variable cost to meet unpredicted requirements. High performance computing capacity is now optimized for predictable (or unpredictable) workloads as well as any project scheduling challenge. Clients can shift their focus and investment to their core business rather than deal with IT complexity, and avoid large upfront capital expenditures associated with their HPC infrastructure. The goal is to gain competitive advantage without the financial burden, management responsibilities, and risks of owning a supercomputer. With DCCoD, customers are able to pursue the optimal combination of in-house resources with high performance computing capacity available on demand. IBM hosts several centers in locations worldwide, equipped with thousands of processors of compute capacity, including Linux clusters, SMP's, as well as access to emerging technologies such as Blue Gene.

DCCoD is being used by IBM's clients and partners to accelerate their time to production, save costs, and optimize business opportunities. One example, in the field of digital content creation, is ***Threshold Digital Research Labs***, a major producer of digital animation and special effects, which was looking for ways to meet production deadlines while minimizing costs. Because the nature of animation and special effects workloads is unpredictable, Threshold could not afford to invest in a huge computing infrastructure just to deal with peak demands. The solution was an on demand e-studio system that provided an unlimited growth path with limited, variable costs. Production time was reduced to half that of typical animation studios, with an expected 25% reduction in animation cost and a 12% to 15% reduction in overall labor costs.

Visualization

Many HPC workloads generate vast quantities of data. Visualization is the key to understanding the full implications of that data. It requires massive computing power and intensive graphics rendering capabilities. Traditionally, visualization has been accomplished by way of specialized and monolithic infrastructures that are costly to acquire, operate, and maintain.

IBM's Deep Computing Visualization (DCV) solution offers an alternative approach. The DCV solution bundles IBM IntelliStation® workstations, commodity graphics adapters, and innovative middleware to deliver an extraordinarily flexible and powerful visualization solution at lower cost.

At the heart of IBM's DCV infrastructure is middleware that links multiple workstations into a scalable configuration, effectively combining the output of multiple graphics adapters. DCV is based on Linux and OpenGL, and delivers two functional modes:

- **Scalable Visual Networking** is a "one-to-few" mode that displays applications on multiple projectors and/or monitors, with the ability to create immersive or stereo visualization environments.
- **Remote Visual Networking** is a "one-to-many" mode that transports rendered frames to multiple remote locations, allowing geographically dispersed users to simultaneously participate in collaborative sessions.

Both Scalable Visual Networking and Remote Visual Networking can be used simultaneously or separately without additional setup.

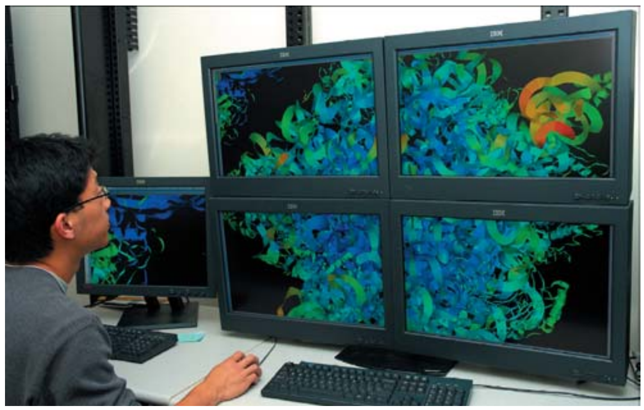

Deep Computing Visualization (DCV) Output

DCV can be used in a wide range of data-intensive application areas, including upstream petroleum, life sciences, and automotive design. A case in point is the ***Carleton University School of Architecture*** in Ottawa, Canada, which hosts the Carleton Immersive Media Studio (CIMS). CIMS promotes the use of IT in architecture design and content delivery.

At CIMS, students participate in advanced studies in visualization and simulation. Using new media technologies, students produce digital content that enhances the development and application of new technologies in a variety of cultural fields. The CIMS research agenda encompasses the development of new technologies and expertise in three-dimensional (3D) real-time visualization and simulation.

IBM's Deep Computing Visualization (DCV) offering was selected as the visualization engine to deliver and display rendering results to remote locations. To address the requirement, CIMS implemented a cluster comprised of IntelliStation A Pro machines with nVidia FX3000G 3D graphics cards, which was used to process rendering instructions from IBM DCV software. CIMS used all of the rendering nodes collectively to contribute to the visualization effect.

The institute chose the IBM DCV solution because a traditional SMP graphic-pipe-based visualization server was too expensive and because the IBM solution enabled CIMS to take advantage of the advanced technology and low price found in commodity graphics cards. The IBM DCV solution also enabled CIMS to scale up performance easily by adding more IntelliStation machines to the cluster, and/or by upgrading the graphics cards.

Special Purpose Systems

IBM often partners with clients to bring new technologies to market. The Accelerated Strategic Computer Initiative (ASCI) program with the US Department of Energy is a prime example of such collaboration. In 1998, ASCI Blue Pacific was installed at Lawrence Livermore National Lab (LLNL), exceeding 3.8 teraflops with IBM POWER technology. In 2000, IBM installed ASCI White at LLNL, achieving 12.3 teraflops using POWER3 technology. The ASCI program was later renamed Advanced Simulation and Computing (ASC), and in 2005 IBM delivered ASC Purple, a p575-based system exceeding 100 teraflops. The ASC Purple contract included the Blue Gene/L system, which was produced in close partnership with LLNL. With the release of the Top500 list in June 2005, the Blue Gene/L machine at LLNL achieved the top position as the world's fastest supercomputer at 136.8 teraflops.

The novel nature of Blue Gene lies in the techniques used to minimize power dissipation and maximize density, and the topology and performance of the interconnection fabrics needed to maintain a fully balanced system. The result is hundreds of teraflops of compute capability, while consuming a fraction of the power and floor space, all at a greatly reduced cost.

Blue Gene is capable of supporting a myriad of applications, but its design point was to accomplish simulations of the most complex phenomena. Blue Gene is uniquely suited for fields of scientific research like astrophysics, genomics, and magnetic fusion. Blue Gene provides the optimal platform for national security simulations, including management of the nuclear weapons stockpile.

Blue Gene also dramatically improves scalability and cost performance for many compute intensive applications, such as biology and life sciences, earth sciences, materials science, physics, gravity, and plasma physics. Numerous Blue Gene systems are now installed in research and academic laboratories worldwide.

Blue Gene/L

Storage

System storage is vital to many Deep Computing applications, ranging from medical imaging to digital content management to data mining and warehousing. Several storage architectures are available from IBM, from on-board to SAN to NAS solutions. Storage virtualization adds capability while simplifying the application's view of storage.

Many Deep Computing solutions include the IBM General Parallel File System (GPFS), a highly available file system for cluster systems. Both Linux and AIX 5L clusters are supported. GPFS supports file systems of several tens of terabytes and runs at I/O rates of several gigabytes per second on clusters of hundreds of nodes. Its low overhead shared disk model allows users and applications on different nodes to share data, while providing the behavior of a local general-purpose file system at each node.

GPFS provides high performance for applications accessing medium to large files through wide data striping, and supports very large file systems. GPFS performance is scalable as processor and disk hardware are added to the system. Moreover, GPFS provides simple and flexible administration of the entire system from a single node. Some of the administrative tasks have parallel extensions for enhanced performance.

In addition to high-speed parallel file access, GPFS provides fault tolerance, including automatic recovery from disk and node failures. Its robust design makes GPFS appropriate for commercial applications such as large web servers, data mining, engineering design, scientific computing, and digital libraries.

The High Performance Storage System (HPSS) addresses the largest data storage needs. HPSS is the result of years of development and collaboration between IBM, five of the US National Labs, and contributions from universities and laboratories worldwide. HPSS provides a flexible and scalable virtual file system for storing the largest amounts of data and accessing it quickly, in a hierarchical fashion. Several HPSS systems currently exceed a petabyte (1024 terabytes) in size, and the architecture is intended to handle exabytes (1024 petabytes).

Under the virtualization layer that gives the appearance of a simple hierarchical file system, the HPSS uses cluster and SAN technology to aggregate large numbers of computers, disks, and tape systems. The HPSS manages the data such that least used data is on tape subsystems, while more frequently accessed data is on disk for faster access. Under the covers, DB2® provides the basis for a scalable metadata engine.

HPSS systems today may provide I/O rates in the range of 100s of terabytes per day, with growth expected to the petabyte range. HPSS provides scalable I/O rates, 24/7 availability, and can recover transactions down to the I/O block level. While the HPSS architecture is standard, each implementation is tailored to the client's particular need. The HPSS is supplied by IBM Global Services.

From a storage hardware perspective, the largest Deep Computing solutions are often implemented on IBM's TotalStorage® DS8000 and DS6000 Fibre Disk. The DS4000 systems are optimal for mid-range systems. IBM's Linear Tape-Open (LTO) Ultrium series tape systems are deployed across a broad range of installations.

- DS8000 and DS6000 are intelligent storage subsystems with built-in storage management functions, such as copying. The DS 8000, the larger of the two, also includes the ability to logically partition the storage, allowing multiple users to share the device with absolute isolation of data. The DS 8000 ranges into the hundreds of terabyte raw capacity, with a design goal of 100 years mean-time-between-failures (MTBF). The embedded processor on the DS 8000 is a POWER5 processor, while the embedded processor for the DS 6000 is a PowerPC®. The DS 6000 is of similar architecture and shares 97% of the code base for embedded functions. Its capacity ranges in tens of terabytes, and has a smaller physical footprint. Both are accessed by multiple 2GB fibre channels and provide the option of highly scalable I/O. Both include user configurable RAID 5 and RAID 10 (mirrored) capability.

- The DS4000 series (formerly FAStT) has been enhanced to complement the entry and enterprise disk system offerings with the DS4000 Storage Manager V9.10, enhanced remote mirror option, DS4100 option for larger capacity configurations, along with support for EXP100 serial ATA expansion units attached to DS4400s.

- Ultrium LTO tape provides the ideal price/performance solution for data archival. The IBM TotalStorage 3580 Model L33 is a third generation LTO device in a stand-alone or rack mounted package. Its capacity ranges to hundreds of gigabytes, with a compression feature that allows an effective doubling of the raw capacity, with maximum access speeds in the high tens of MB/s.

Workstations

Far more powerful than a PC, workstations combine the fastest processors with high function graphics. Most contain two CPUs, large memory and storage, and state of the art graphics adapters (with large graphics memory). IBM's IntelliStation delivers optimum performance, with a choice of architectures and chip sets. IBM clients are using the IntelliStation series in diverse fields such as engineering, digital content creation, product design, movie production, chemical and petroleum analysis, and life sciences.

IBM IntelliStation Pro workstations are designed for the most challenging business operations. IntelliStation Pro workstations are available with single or dual Intel or AMD processors, and run Windows® or Linux. With very large memory, terabytes of disk space, and excellent 3D graphics performance, the IntelliStation Pro series offers 32- and 64-bit performance for the most demanding engineering and graphics needs.

IntelliStation POWER is an innovative UNIX workstation for power users. The IntelliStation POWER features many of the mainframe-inspired technology and autonomic computing reliability, availability and serviceability (RAS) capabilities available in IBM UNIX operating system-based servers. The workstation includes the AIX 5L operating system with its extensive security, workload management efficiencies and excellent 3D graphics performance. The IntelliStation POWER workstation is packaged as a mid- to high-end deskside workstation designed to provide the power and RAS capabilities needed for MCAD, front-end graphics processing or other floating-point-intensive applications. Like the Pro, its features include very large memory and terabytes of disk space, with excellent 3D graphics performance to support the most demanding engineering and graphics tasks.

Grid Computing

Grid computing enables the virtualization of distributed computing and data resources such as processing, network bandwidth and storage capacity to create a single system image, granting users and applications seamless access to vast IT capabilities. Just as an Internet user views a unified instance of content via the Web, a grid user essentially sees a single, large virtual computer.

At its core, grid computing is based on an open set of standards and protocols (e.g., Open Grid Services Architecture) that enable communication across heterogeneous, geographically dispersed environments. With grid computing, organizations can optimize computing and data resources, pool them for large capacity workloads, share them across networks and enable collaboration.

Grid architecture underlies Deep Computing Capacity on Demand. At the other end of the spectrum, IBM recently announced a Grid and Grow offering that uses BladeCenter hardware, a choice of operating systems, and grid enabled software, preconfigured and assembled. It combines the convenience of a turn-key system with the flexibility of as-needed growth.

Grid computing is frequently used as a means of virtualizing HPC resources and data to accelerate overall performance and throughput. One such example is the US ***National Digital Mammography Archive (NMDA)*** hosted at the University of Pennsylvania, which developed a revolutionary Electronic Medical Record (EMR) data grid and repository. The repository provided a framework for a patient-centric medical record system that could capture, from any location, the full range of healthcare files including digital images, records, and clinical history. The goal was to provide access to current and past patient records in two to 90 seconds, ensuring optimized record review for diagnoses while maintaining the utmost security and privacy requirements. The solution was comprised of IBM eServer systems, DB2 Universal Database™ and GPFS. The three-tier architecture of the NDMA grid leveraged the strengths of the IBM eServer systems with open protocols from Globus®. The use of AIX 5L, Linux and Windows illustrated the inherent heterogeneity of computing grids.

System Software and Applications

Many clients have multiple system architectures, and it is critical for those systems to be integrated into a seamless operating unit. IBM's HPC software suite enables this integration by allowing clusters to be managed from a common, centralized management tool, by allowing multiple systems to access shared file systems under GPFS, by allowing jobs to be scheduled across the entire system, and by enabling grid access to the system. Applications further benefit from common compilers and run-time libraries, reducing the porting and support effort required.

IBM's Deep Computing platforms include operating systems support for AIX 5L and Linux, as well as Windows. AIX 5L is IBM's premier platform for UNIX applications, and is optimized for POWER hardware. IBM is fully committed to the open source movement and supports Linux as a key platform for computing in the 21st century. IBM will continue to work with the open source community, bringing relevant technologies and experience to enhance Linux, to help define the standards, and to extend Linux to support enterprise wide systems.

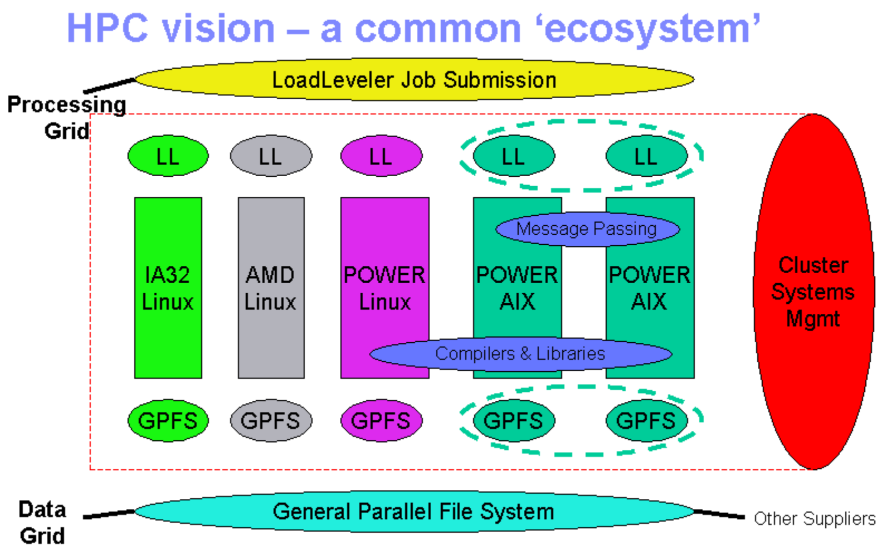

The system and storage management portfolio includes IBM Cluster Systems Management (CSM) on the Cluster 1350, CSM and Parallel Systems Support Programs (PSSP) on the Cluster 1600, and IBM General Parallel File System (GPFS) for faster I/O on both platforms. IBM provides compilers and libraries for IBM's UNIX systems, including the AIX 5L compiler suite, recognized as one of the best optimizing compilers available, the Engineering Scientific Subroutine Library (ESSL), and a parallel version of ESSL (PESSL) for computational efficiency. IBM LoadLeveler® provides sophisticated job scheduling. For Cluster 1350 middleware, open source and commercial schedulers are also available.

Complementing IBM’s robust HPC software stack, IBM forms partnerships with many Independent Software Vendors (ISVs) to offer a broad portfolio of solutions, spanning across all of the Deep Computing markets – life sciences, petroleum, PLM, digital media, business intelligence, and more. This HPC software portfolio is complemented by a wide range of commonly used open source and public domain tools, libraries and other packages.

Services

Through IBM Global Services, IBM offers services for Deep Computing ranging from installation of complex systems to the complete operation and management of the data center. Many Deep Computing solutions are packaged by IBM Global Services, and IBM Global Finance can also provide favorable terms for financing.

Using Linux Clusters as an example, IBM can provide services for installation, customization, and migration from Unix to Linux. IBM Global Services can offer training for the administration of the environment as needed. Offerings can be tailored to be very application specific, using expertise and consulting from organizations ranging from IBM Research to the IBM Computational Biology Center.

Systems Positioning

IBM approaches Deep Computing not as a monolithic market best served by a single solution, but as a market of highly varied applications which require different types of systems and solutions. One size does not fit all.

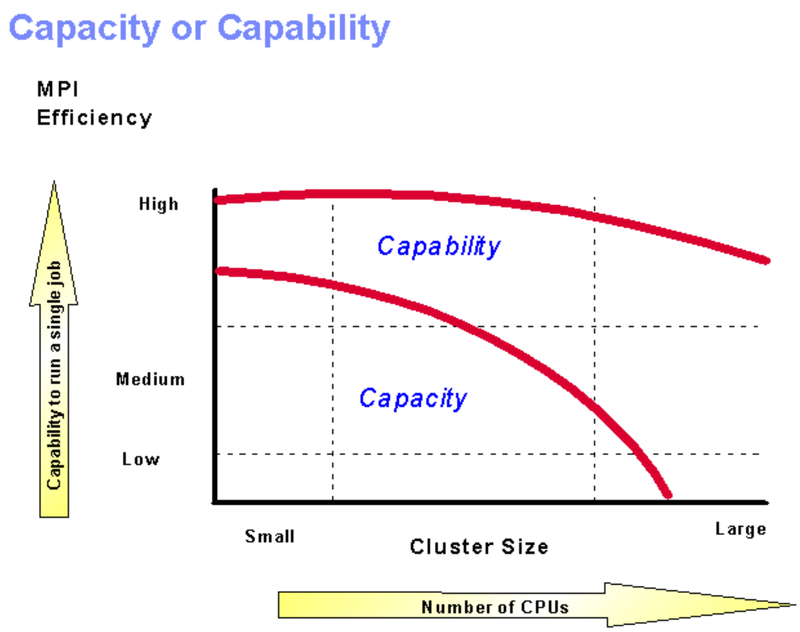

Deep Computing systems can be characterized as being either *Capability* or *Capacity* machines, using IDC®'s terminology. *Capability* systems are able to run a single job using all, or most of, the computing power of the whole machine. As the job is parallelized across many processors, intercommunication between the processors is often high, and efficient execution can only be achieved if the message passing (commonly using MPI) efficiency is high. *Capacity* systems are used to run multiple smaller jobs (often multiple instance of the same job), none of which extend to run across the whole machine, and can therefore be implemented with standard interconnects while maintaining high performance.

IDC characterizes capacity systems across a range of price bands, from larger enterprise and divisional clusters and SMP's, to smaller departmental and workgroup systems - typically single SMP servers or small clusters. Departmental clusters most often run Linux and may use Gigabit Ethernet, Myrinet, or InfiniBand as interconnects.

Both capability and capacity jobs may be run on either "scale-up" or "scale-out" machines. Scale-up servers are usually wide SMP's with many processors; however, many of those SMP servers may be clustered together to achieve the required performance, giving rise to the term scale-out.

HPC – scale-up / scale-out

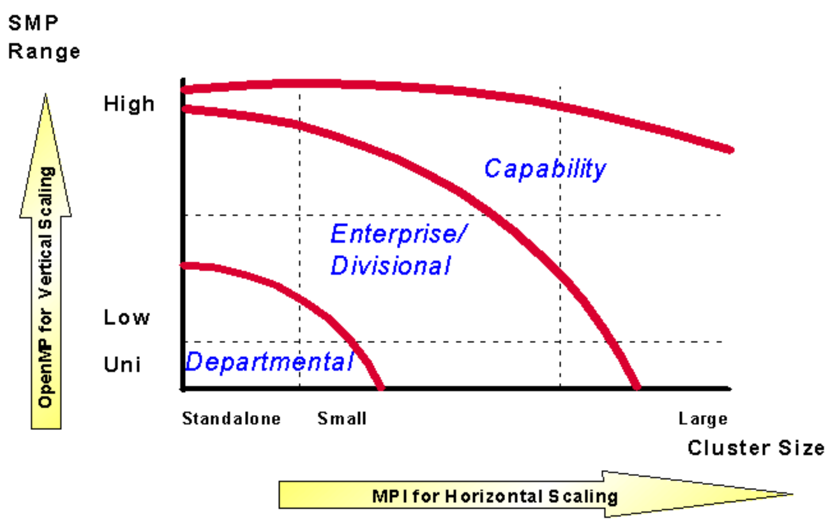

Scale-up (SMP): SMP's are the simplest programming model, with strong ISV support. The operating system may be Linux or UNIX. SMP's are viewed as simple to administer and use, and are often used as throughput systems. Typical applications include Noise, Vibration and Harshness (NVH) in the automotive and aerospace industries, and some applications in bioinformatics and petroleum.

Scale-out (Commodity): These are typical COTS clusters. Performance is "good enough" for many requirements at an excellent price point. Applications include bioinformatics and selected proteomics, certain quantum chemistry and molecular modeling codes, some Computational Fluid Dynamics (CFD) codes, seismic, EDA, and financial modeling.

Scale-out (High Value): These systems are characterized by very high processor power, memory size, and I/O and memory bandwidth. The systems may have thousands of nodes, and the nodes may be larger SMP's. The MPI stack tends to be more mature than the commodity systems. Applications include atmosphere and climate modeling, crash codes for automotive and aerospace, CFD, certain quantum chemistry and molecular modeling codes, reservoir modeling, as well as data mining and data analytics. Blue Gene is an example of the high-value scale-out class of machine. In one system, with one system image, it can have over one hundred thousand processors.

IBM Deep Computing Solutions

There is a vast difference between producing a product and engineering a solution. Solutions entail choosing the optimal product building blocks, both hardware and software, combined with services where appropriate, and integrating them to address specific client and industry needs. A properly designed solution combines subject matter expertise of the application area with a broad knowledge of systems and architectures. IBM's breadth of application expertise, spanning virtually all HPC disciplines, coupled with its leadership product portfolio and choice of architectures, make it uniquely qualified as a solution vendor. IBM brings the full toolkit, not just the hammer.

Solutions are tailored to the client's needs, but generally have common components within a given application area. These components are well integrated and balanced to create a flexible, scalable, and responsive high performance computing infrastructure. IBM offers a range of integrated solutions in Digital Media, Life Sciences, Business Intelligence, and other areas, as exemplified in the sections which follow.

Digital Media

With the virtual explosion of digital data and content, comes an ever-increasing need for systems to process, store, analyze and manage this data. IBM's portfolio for Digital Media solutions includes digital content creation, management and distribution, online games, security and surveillance, video on demand, wireless and mobile content services, and more.

One such solution is the Digital Media Center (DMC), which provides a shared, highly-scalable storage pool, supporting a variety of leading digital broadcast production toolsets and the content created or processed by them. The solution combines integration services with leading IBM technologies, including Tivoli® Storage Manager and IBM TotalStorage LTO Ultrium (Linear Tape Open), GPFS, and pSeries® and xSeries® hardware. With open standards and interoperability as key design points, the Digital Media Center integrates quickly into customers' broadcasting infrastructure.

The ***Australian Broadcast Company (ABC)*** is one of many customers who have successfully implemented DMC. With help from IBM's Global Services, ABC has embarked on a project to digitally archive and store all of the ABC's television and radio history, preserving the programs for years to come. In an Australian first, more than 40 years of television and 60 years of radio history stored on over 100,000 aging analogue tapes will be converted into 500 terabytes of digital data. Each program will then be immediately available online via the ABC's digital network, allowing staff to instantly access, re-use and adapt the programs as required. Converting the analogue tapes will also prevent sound and quality degradation characteristic of analogue magnetic media, preserving one of Australia's most important broadcast archives.

Another Digital Media solution is the IBM Realtime, Adaptive, Multi-Intelligence, Multi-Media Platform (RAMMP). Aimed at the defense, intelligence, security and law enforcement communities, RAMMP enables users to rapidly collect vast amounts of audio, video, image and sensor-based data into a centralized and secure repository, which can then be processed and analyzed to generate real-time results. This solution combines IBM eServer systems with open-standards based software including IBM DB2 Content Manager, IBM DB2, IBM Websphere, and GPFS, together with business partner applications and services.

Healthcare and Life Sciences

The Healthcare and Life Sciences industry addresses a wide range of critical needs for hospitals and clinics, researchers, biotechnology and pharmaceutical companies, and the vast network of providers, insurers, and others who serve this community. As in other areas of Deep Computing, there is a need for a high performance computing infrastructure to conduct research, develop new drugs, collaborate among various institutions, address regulatory compliance issues, and handle the explosion of medical and clinical information which is being generated in a digital form.

From Linux clusters and blade-based solutions which are well-suited to research and development, to the ultra-scale capabilities of Blue Gene for the most demanding life sciences applications, to medical imaging and archiving solutions, IBM offers a broad range of capabilities for this high growth industry. One such example is the IBM eServer BladeCenter Solution for Bioinformatics. Combining IBM's leading JS20 blade technology with Linux and commonly used bioinformatics applications, this highly scalable 64-bit platform enables researchers to analyze complex scientific data more quickly and at lower cost.

The IBM Medical Archive Solution (MAS) is a pre-configured, hierarchical storage solution designed to meet the online, near-line and off-line image management requirements of the healthcare enterprise. It enables timely distribution of exam results and timely access for radiologists and clinical staff, while meeting or exceeding government regulations (e.g. HIPAA) for security and privacy. This solution features IBM Tivoli and GPFS software and TotalStorage technologies to manage and access medical images in a secure and cost-effective manner. IBM recently introduced the Grid Medical Archive Solution (GMAS), offering an even more flexible and scalable storage solution for healthcare organizations and hospital networks.

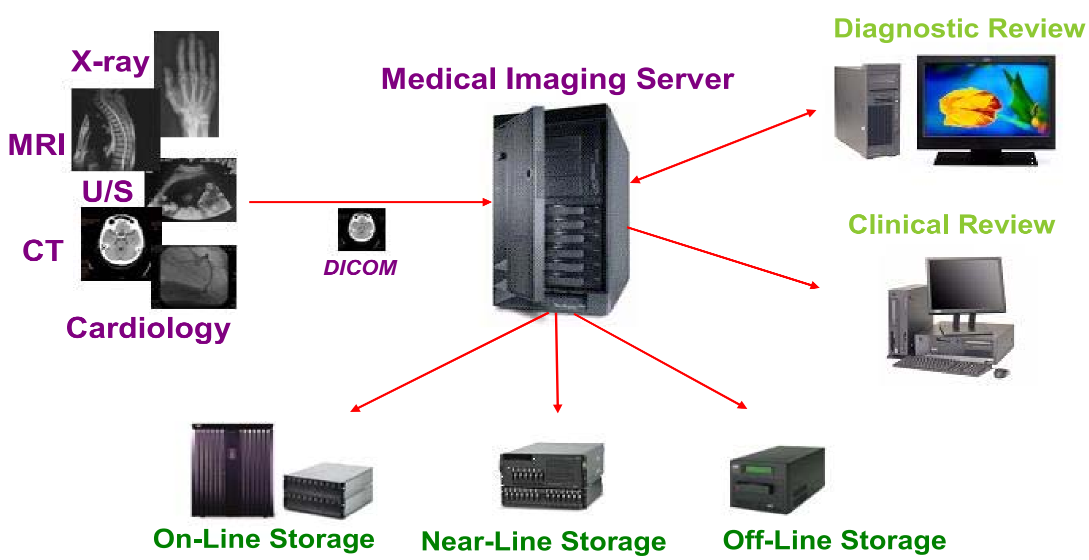

Institutions around the world are employing IBM's life sciences solutions to facilitate and accelerate the pace of their research. The mission of the ***German Cancer Research Center (DFKZ)*** in Heidelberg, Germany, is to make fundamental contributions to the medical community's understanding of the mechanisms of cancer development and cancer risk factors. The results of its research are intended to lead to new approaches to cancer prevention, diagnosis and therapy. The DFKZ employs standard computational chemistry applications such as Gaussian, Amber, and TurboMol, to simulate the reactions that take place between molecules. All are floating-point-intensive applications which analyze large data sets. To address the high performance computing requirements at DFKZ, IBM delivered an optimal solution comprised of IBM pSeries servers and Total Storage. Two POWER4 p690 systems with a total of 96 CPUs provided 330 billion floating-point operations per second, and IBM TotalStorage Linear Tape-Open (LTO) tape libraries were deployed in two locations. Together, they could process and analyze approximately 1TB of data at any given time.

The ***University of Alberta*** is one of Canada's foremost research-intensive universities. It houses the Institute for Biomolecular Design, an interdisciplinary research and platform technology center. The institute conducts research with the aim of simulating all biological processes that take place within a single cell. To simulate a biologically and chemically accurate computer model of a living cell, the institute needed to gather and analyze a huge amount of data and perform significant application development effort, requiring access to a high performance computing environment. The Institute for Biomolecular Design engaged IBM to deploy their HPC infrastructure on a foundation of IBM hardware, including a pair of pSeries 690 servers running AIX 5L V5.2, connected to an IBM TotalStorage DS4500 with 6TB of storage capacity, with GPFS as the primary file system.

Business Intelligence

As companies or organizations seek to gain better insight about their clients, their competitors, or shifting market conditions, it is critical that they have the tools to ingest and analyze vast amounts of data and produce real-time results. Retailers, banks, governments, and other institutions are employing sophisticated data warehousing and data mining techniques to deal with this explosion of information and gain competitive advantage. IBM offers Business Intelligence solutions and consulting services to help clients navigate through this maze of data.

One such solution is the IBM Data Warehousing Balanced Configuration Unit for AIX (BCU). This end-to-end solution is comprised of IBM eServer p5 575 and IBM TotalStorage DS4500 hardware with the IBM DB2 Data Warehouse Edition, integrated into a scalable package that offers optimal price/.performance for data mart consolidation and data warehousing. The BCU offers a simple yet scalable design to handle growing data needs.

The breadth of Deep Computing solutions goes well beyond the examples described here. These end-to-end solutions draw upon IBM's robust hardware and software portfolio, combined with products and technologies from business partners and ISV's, and integrated with the deep expertise IBM delivers through its services and consulting organizations. IBM has a proven track record of developing, integrating, and delivering complete high performance computing solutions to address very complex and difficult problems. That requires innovative technologies, breakthrough initiatives, and thought leadership. This solution-oriented approach is what sets IBM apart in the industry.

Conclusion

IBM's leadership in Deep Computing is evident across many dimensions: a long-standing track record of innovation and investment in high performance computing, a passion for excellence, the strength and diversity of the product portfolio, and the breadth of solutions aimed at multiple industry areas. IBM has unrivalled expertise in architecting and delivering HPC systems across a wide spectrum, from simple, low-end installations to the largest supercomputers addressing the "grand challenges" in science, engineering and research.

IBM has been investing in high performance computing for almost half a century. No other vendor comes close. IBM Research, perhaps the pre-eminent corporate technical research group in the world, invented RISC microprocessors and relational databases. Research parallel machines paved the way for industrial strength real world parallel machines, and innovation in switch technologies enabled much higher degrees of scalability. Along the way, Deep Blue became the world chess champion. More recently, the Blue Gene supercomputer, based on the Power architecture, has demonstrated unparalleled performance with an order of magnitude reduction in floor space and power consumption.

IBM's leadership is also evidenced by its position on the Top500 Supercomputer List (http://www.top500.org), which tabulates the 500 largest supercomputers worldwide. As of June 2005, IBM led the list with the fastest supercomputer in the world (Blue Gene), the highest installed aggregate throughput, more entries than all other vendors combined as well as the most systems in the top 10, 20 and 100, and the greatest numbers of Linux clusters. IBM systems on the Top500 list are being used to run some of the world's most advanced scientific and engineering applications, including energy and environmental science, weather and climate research, biosciences, computational chemistry, electronic design automation, and materials sciences.

If raw compute power were the only criteria for high performance computing, then many vendors could claim leadership. It's not that simple. First, innovative and disruptive technologies require years of research and development. Second, making those technologies useful in the real world requires deep technical know-how and extensive industry knowledge. Finally, designing industry-specific high performance computing solutions to exploit new technologies is a very complex issue.

IBM views Deep Computing as a vital opportunity, and a market that IBM is uniquely positioned to serve. Leadership in this market is key to leadership in the other markets that IBM serves, since technologies developed here will inevitably make their way into the broader commercial marketplace.

Far from being a "box" vendor, IBM understands that the value it can bring to the market is its breadth of innovative solutions and expertise in system architectures, compilers, tools, middleware and applications. From Fast Fourier Transforms to financial modeling to financing your system, IBM offers the complete solution for Deep Computing.

Authors:

NormSnyder LLC
Norm Snyder, President
NormSnyder@ieee.org

Rebecca Austen, IBM
Dinh Pham, IBM

The IBM home page on the Internet can be found at: http://www.ibm.com.

The IBM Deep Computing home page on the Internet can be found at: http://www.ibm.com/deepcomputing.

IBM Corporation
Systems and Technology Group
Route 100
Somers, New York 10589

Produced in the United States of America
November 2005