Cabot Partners Group, Inc. 100 Woodcrest Lane, Danbury CT 06810, www.cabotpartners.com

Cabot Partners
Optimizing Business Value

# Beyond PetaFlops: Scalable, Energy Efficient IBM System x iDataPlex dx360 M4 powered by Intel Xeon processor E5-2600 Product Family

**Srini Chari, Ph.D., MBA**
**March 2012**
chari@cabotpartners.com

***Sponsored by IBM***

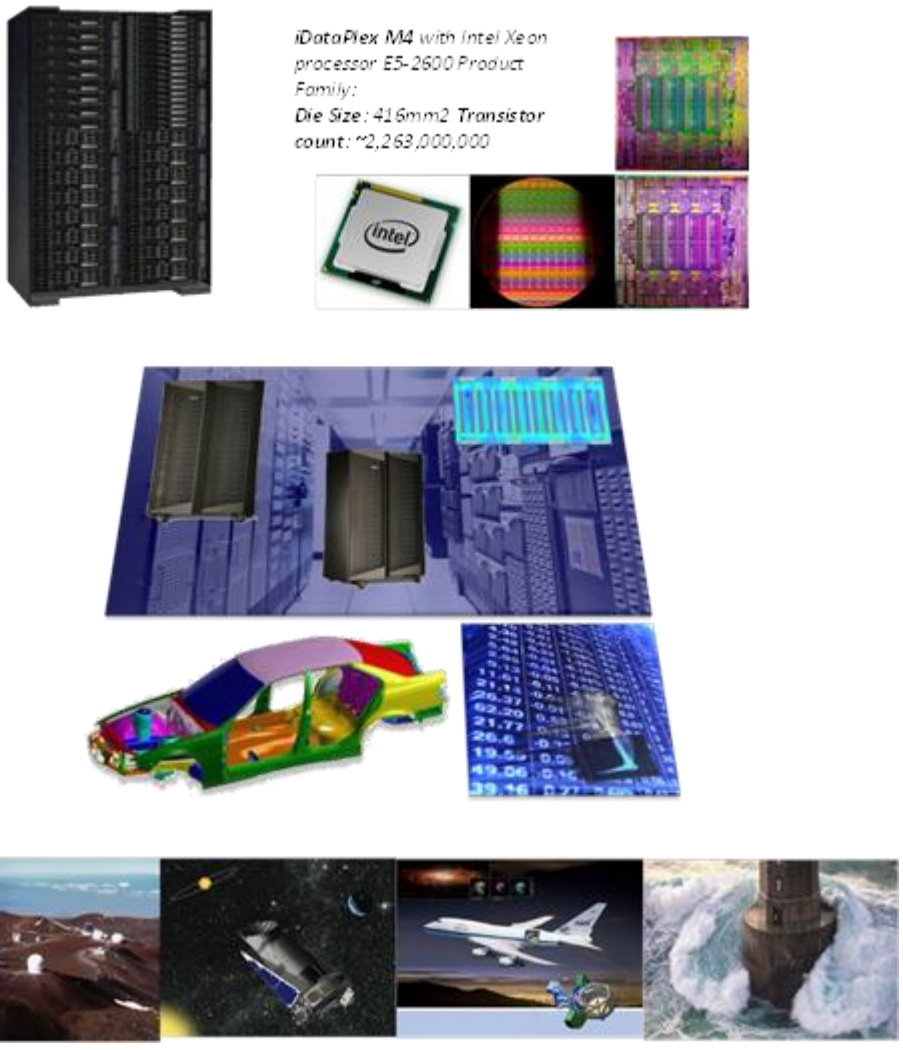

# 1. Introduction

Enterprises today increasingly face challenges in order to keep pace with growth and varying business demands that require rapid scaling and technology evolution in the data centers. While each technology refresh cycle delivers a significant performance boost at roughly the same price, enterprises continue to see rising facilities costs and mounting energy bills for power and cooling. Other key components of the data center Total Cost of Ownership (TCO) include operational costs for labor and ongoing equipment maintenance costs. To get the most out of their current and future IT investments, data center managers must minimize TCO while achieving desired efficiencies, computing performance and effective manageability of massive scale-out environments.

As a leading provider of High Performance Computing (HPC) solutions across a wide range of application and industry segments over the last few decades, IBM has constantly innovated and helped alleviate customer pain points through its rich portfolio for HPC solutions. The new system design of IBM's System x iDataPlex dx360 M4 servers leverages Intel's industry-standard highly scalable and powerful processors to address energy efficiency and significantly reduce carbon emission footprints of next generation data centers. Intel's 32nm Xeon E5-2600 product family processor powers iDataPlex dx360 M4 air or water-cooled half depth servers to provide more rack density per square foot of floor space in the data center, thus enabling higher compute power with lower energy consumption.

With this latest iDataPlex offering, IBM extends its current HPC cluster offerings by providing an economical architecture that can increase computing density in a standard enterprise rack and potentially scale-out to ExaFLOPs with optimum energy and space footprint. In this paper, we list some of the challenges faced by Petascale HPC clusters and other large data centers and show how IBM is addressing these challenges effectively through its innovative and energy efficient dx360 M4 servers.

## Overcoming Data Center Challenges

Business users, engineers and scientists attempting to solve challenging problems in engineering, financial markets, media, public sector, oil & gas, life and earth sciences rely heavily on HPC systems. These users are some of the most demanding IT solution clients. Their insatiable demand for computational performance continues to drive HPC platforms towards increasing system scalability, larger computing density, lower power consumption, efficient system cooling, easier and more reliable system management, and adaptable programming environments.
Multiple industry studies indicate that revenue from HPC servers will continue to grow much faster than overall server revenues with clusters being the dominant platform – over 75 percent of HPC servers[1]. The battle for HPC leadership is becoming increasingly global and strategic. Soon, more real-world applications will run at trans-petaflop speeds. However, these studies also suggest that the escalating costs of server management, power and cooling, reliability, availability and serviceability (RAS)[2]of HPC clusters with thousands of racks and facilities management at HPC data centers will far outpace the costs of buying new servers. This has caused a severe crisis in HPC data centers. IT solution providers such as IBM have responded to this by designing innovative solutions that address these issues while retaining all the attractive attributes of industry-standard cluster architectures for HPC.

[1] HPC Market figures, trends and predictions by IDC at ISC, June 2011 http://www.hpcwire.com/hpcwire/2011-06-21/idc_shares_hpc_market_figures,_trends,_predictions_at_isc.html
[2] Cabot- Partners' HPC RAS paper http://www-03.ibm.com/systems/resources/systems_deepcomputing_IBMPower-HPC-RAS_Final-1.pdf

## HPC Cluster Differentiators: iDataPlex Innovations in Packaging, Power and Cooling

IBM's iDataPlex dx360 M4 server can significantly reduce data center costs arising from power consumption. It also helps reduce floor space requirements for cold aisles through innovations in cooling technologies and smart system designs that require less floor space. Data centers can deploy the next generation iDataPlex dx360 M4 servers and obtain more computing density while consuming the less power. They also benefit from new packaging and cooling methods - traditionally used in very large enterprise class systems and proven very effective for designing energy-efficient clusters where water is pumped through the rack and into node-level cooling loops, which will provide data centers with un-paralleled advantages.

The dx360 M4 server is available as a hybrid solution – both with air and direct water-cooled design. The direct water-cooled solution is innovative node level cooling that pumps water through the rack and into node-level warm-water cooling loops. This water-cooled design eliminates the need for fans in a 2U chassis and significantly boosts energy efficiency. This allows operations that are more reliable and better placement of large clusters at HPC data centers - significantly reducing the overall power, cooling and facilities management costs for the entire data center.


**Figure 1: Cooling innovation in IBM's iDataPlex dx360 M4 (Source IBM)**

## iDataPlex: Scalable, Energy Efficient and Economical HPC Workhorse

With direct water-cooling technology, iDataPlex can help lower power consumption by up to 40% as compared to air-cooled iDataPlex systems[3] and can greatly reduce the air conditioning need in data centers. Another advantage is reuse of 90% of the heat removed using direct water-cooling technology from the data center for community purposes such as heating buildings or water. The direct effect of this innovative architecture is that it economically increases the compute density in the data center by a factor of five while retaining many of the attractive attributes of current

[3] Air-cooled iDataPlex systems themselves are 40% energy efficient than older air-cooled systems. The water-cooled iDataPlex is 40% more energy efficient than latest air-cooled iDataPlex. (Source IBM)

integrated rack cluster solutions from IBM. As compared to air-cooling, direct water-cooling technology also consumes much less energy.

### IBM Continues Active Collaboration with HPC Developers and Partners

Since iDataPlex can provide a large scale-out cluster using industry standard components, the large portfolio of HPC applications that have been mapped, migrated, and optimized for the IBM Cluster architecture will benefit greatly from this increased density, scalability, performance and cost-effectiveness. This helps to protect customer's prior investment in application enablement and related. iDataPlex server provides an efficient, reliable and cost-effective building block for supercomputing HPC clusters.

## 2. HPC Data Centers are in a Crisis

The increased use of HPC in mission critical business processes across multiple industries, the growing analytical complexity of interdisciplinary analysis, and the explosive growth of data that

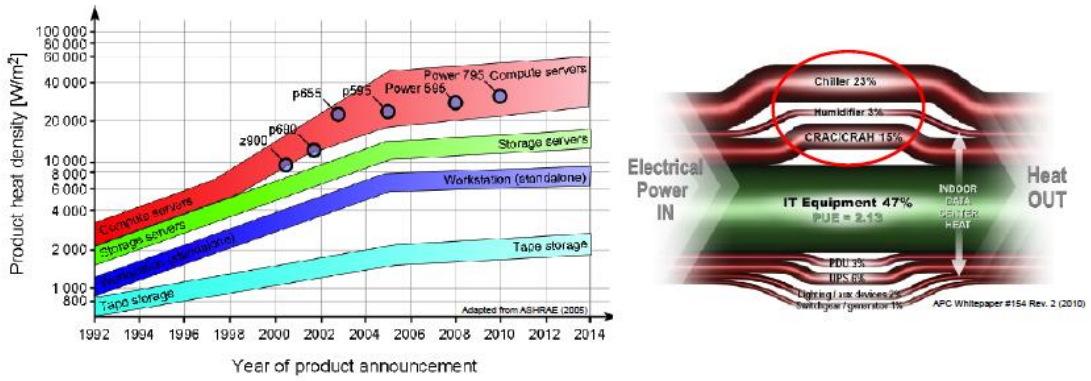


**Figure 2: The Energy challenge for Data Centers: Cooling (*Source: IBM Zurich*)**

need to be analyzed have contributed to the spectacular use of scalable, high density computing on large multi-core clusters. However, the rising costs for power and cooling, systems management, RAS and facilities management have resulted in HPC data center crises.

### Data Explosion and Growing Analytical Complexity Have Precipitated a HPC Data Center Crisis

Data volume is growing exponentially in almost every HPC discipline. The increasing diversity of the available data further complicates the process of extracting valuable insights from gigantic volumes of data. Synthesizing data from disparate sources in order to leverage insights from a holistic point of view of global events or business processes has become a fundamental necessity for competitive advantage for most enterprises that rely on HPC. But Big Data analytics and synthesizing, analyzing digital forms of technical or business data such as e-mail, instant messages, blogs, Web sites, data from electronic sensors, diagnostic images, media from television, radio, and audio and video – cannot be achieved efficiently by the current approaches that today's systems were built to support. The increased complexity of the statistical and mathematical models used for solving inter-disciplinary engineering and scientific problems is further fueling the need for greater HPC capability.

At the same time, merely adding more capacity in terms of servers and increasing compute density does not solve the problem either. Instead, it makes it more complex and unviable with rising power, energy and facilities costs and trade-offs involved in reliability, availability and serviceability issues.

Today, the data center infrastructure has reached the breaking point in terms of energy and cooling because of the massive explosion of data volume and the associated compute resources and floor space required in order to analyze this data. Studies indicate that close to one-third of server power and cooling issues affect purchases[4]. Some enterprises cannot find space or electricity to install new servers, which sit in closets waiting for new construction. Enterprises are looking at solutions that use the minimum possible power and other services in order to be able to accommodate expansion plans. The winning combination that can survive this data and analytics deluge would include high density computing with low energy consumption throughout the business process, end to end.

## New Environmental Mandates plus Rising Energy Costs and Demand Have Exacerbated HPC Data Center Crises

The consumption of energy in large data centers for power and cooling is likely to increase steadily during the next decade. In near-term, unless there is a breakthrough and development of more economical approaches to producing energy, the energy costs are likely to increase even more rapidly. Moreover, new environmental compliance and governance mandates have put requirements that are even more challenging on IT infrastructure solution providers. In addition to existing metrics for performance and price/performance, the IT industry is defining additional metrics such as gigaflops/watt that rate HPC solution providers. The Top Green 500 list[5] is becoming as important as the Top 500[6] list of supercomputers as HPC solution providers compete for "bragging rights". Recent studies[7] on data center energy distribution show increasing heat density in various IT components, particularly in servers, storage, and networking. With 60% of the energy[8], going towards cooling and power, more effective power and cooling solutions are required in an HPC system and in the entire HPC data center. Going a step further, IBM is looking at addressing energy problems such that the data center is no longer an energy sink but becomes an important entity in the energy landscape through energy reuse. Innovations and technologies such as the 'direct water cooling' in iDataPlex indicate IBM's novel approach to multi-functional data centers[9] with evolve from being a pure energy consumer to active participants in a distributed and interconnected energy grid.

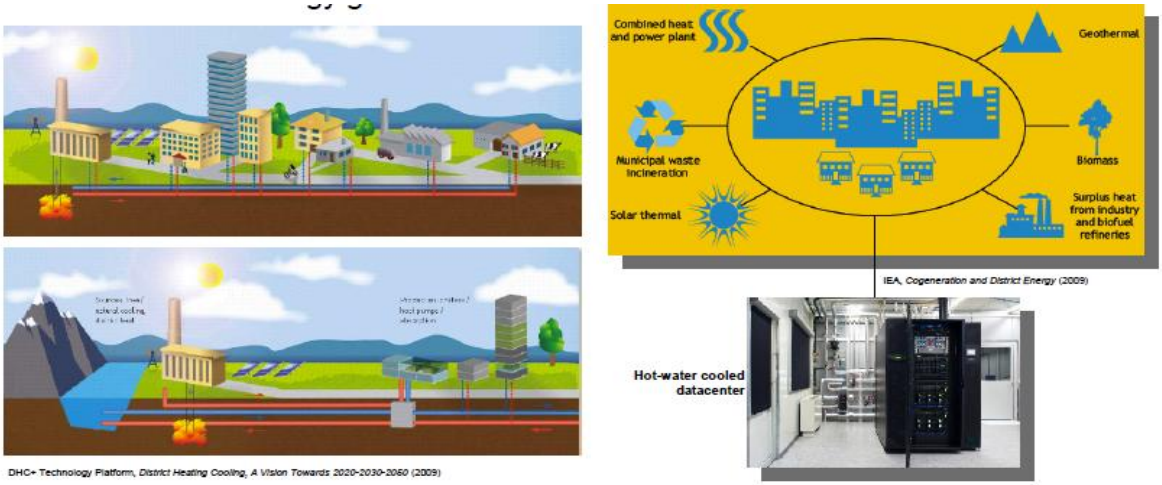


**Figure 3: Future: Multifunctional Data Centers (Source IBM Zurich)**

[4] IBM Whitepaper - Into the Future: The Fast Track for Financial Markets https://www-304.ibm.com/easyaccess/fileserve?contentid=121196
[5] The Green 500 List, www.green500.org/Lists.html
[6] The Top 500 list at www.top500.org
[7] http://www.upsite.com/TUIpages/whitepapers/tuiheat1.0.html
[8] Data source: Creating Energy-Efficient Data Centers, , U.S. Department of Energy , Data Center Facilities and Engineering Conference , May 18, 2007
[9] Recycling Thermal Energies boosts efficiencies in Data Centers and concentrated photovoltaic systems, IBM Zurich http://www.esc.ethz.ch/ses11/Ruch_presentation.pdf

## 3. iDataPlex: For High Performance and a Flexible Data Center

To address the new data center challenges, a new way of designing data centers and the server infrastructure that goes into the data centers is essential. In addition to the performance, cost, and scalability requirements at the server level, the data center design must encompass other costs such as those for power, cooling, management, capital acquisition costs as well as the operating costs as the chief design goals. While IBM has years of experience designing server technologies for scale-up and scale-out that primarily focus on performance and scalability as the fundamental requirements, iDataPlex solution focuses on a different set of goals:

- Reduce the initial hardware acquisition costs and on-going maintenance costs for data center owners
- Improve efficiency in power consumption
- Eliminate data center cooling requirements
- Achieve higher server density within the same footprint as the traditional rack layout
- Simplify manageability for massive scale-out environments
- Reduce the time to deployment through pre-configuration and full integration at manufacturing

As is evident, these design goals go far beyond a single server or a single rack level; they are goals for the entire data center. With this new philosophy and the new design, IBM's iDataPlex solution promises to address the data center challenges at various levels:

- an innovative rack design for higher node density
- a flex node chassis based on industry standard components
- shared power and cooling components to improve node as well as rack level efficiencies
- optional rear-door heat exchangers that virtually eliminates traditional cooling and lastly,
- optional direct water cooling technology that can lower power consumption by up to 40%

The following figure highlights evolution of cooling technologies in IBM's iDataPlex solution.

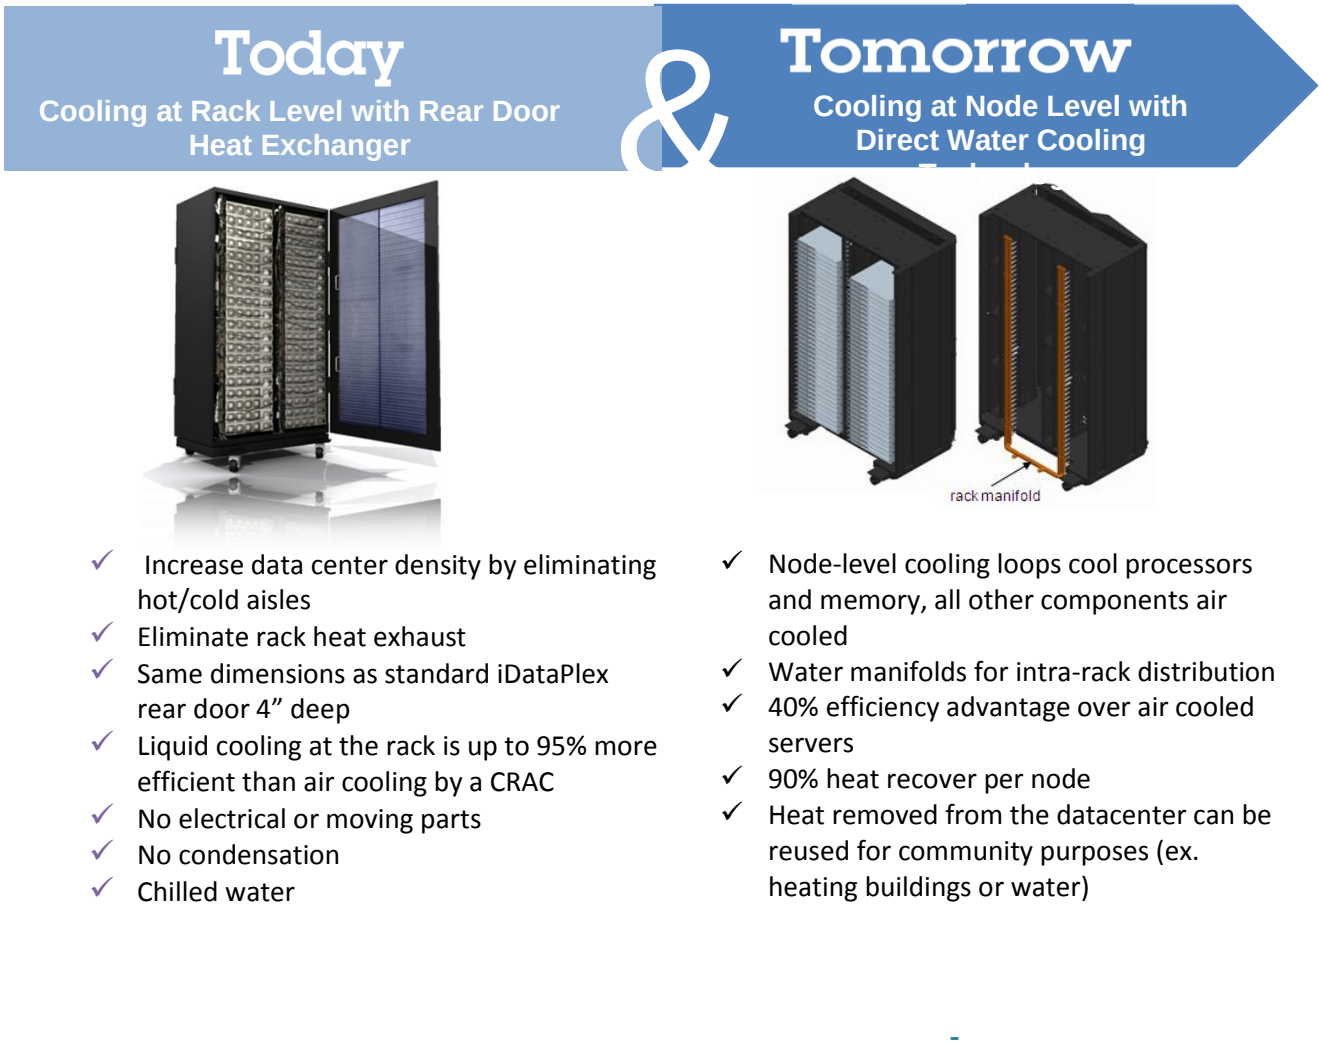


**Figure 4: Water Cooling Evolution of iDataPlex technology (Source IBM)**

## iDataPlex Architecture: Cost Effective, Flexible and Scalable

With major cooling innovations, the direct water-cooled iDataPlex promises to be a unique offering for HPC data centers. IBM's enterprise experience and traditional HPC needs combine to create a low cost platform with no redundant hardware, easily available industry standard parts, extreme compute density, increased performance/watt, and quick and large-scale deployment. With a more efficient standard form factor, iDataPlex combines servers, chassis, switches, PDUs, management appliances, and efficient rear door heat exchanger and latest direct water-cooling technology to address many HPC needs such as flexible configurations for nodes, storage and I/O capability including power-optimized architectures.

With the Intel Xeon E5-2600 series processor powered iDataPlex M4, Petascale HPC clusters and data centers can achieve over 55% reduction in overall TCO and as much as 90% savings[10] in power and cooling through innovations in air cooling, rear door heat exchanger chassis, energy efficient Intel processors, rack, and optional liquid cooling in an iDataPlex. There is an optional iDataPlex Rear Door Heat eXchanger (RDHx), which uses a sealed rack door containing chilled water to cool the rack. The warm-water cooling solution removes up to 90% of the heat from the rack; the RDHx can remove up to 115% of the heat (See Figure 4), actually helping to cool the rest of the data center. Thus, iDataPlex based cluster solutions can potentially reduce the need for CRAC units and free up floor space for additional racks of equipment.

## iDataPlex Features and Benefits at a Glance

iDataPlex is a result of a decade of system innovation by IBM and promises to close the loop between low-end inexpensive volume servers and enterprise class SMP and blade servers. This massive, scalable computing workhorse includes a wide-range of flexible options that make it a great choice for HPC applications and other large-scale applications. The figure below depicts the salient system details and summarizes iDataPlex features and benefits.

| Benefits | iDataPlex Features |
|---|---|
| **Deployable** | ▪ Racks fully integrated by IBM manufacturing at the factory site |
| | ▪ Delivered, installed, and powered up at the data center before customer acceptance |
| | ▪ Minimal on-site time required - shipping dock to image-load |
| **Serviceable** | ▪ All front access - eliminates need for accessing rear of rack |
| | ▪ Blade like design with chassis docking into power connector |
| | ▪ Innovative cable layout provides highly efficient cable routing |
| | ▪ Flexible support options from self maintenance to 24x7x4 response time |
| **Efficient** | ▪ Lower power consumption by up to 40% per rack[11] |
| | ▪ PUE ~1.1 (SuperMUC, Dr. Bruno Michel, IBM Research – Zurich, Manager Advanced Thermal Packaging June 16, 2011) |
| | ▪ Node-level direct water-cooling can help reuse 90% extracted heat for community purposes (heating buildings) |
| | ▪ Rack level liquid cooling is up to 95% more efficient than air cooling by a CRAC |
| | ▪ Energy expenses reduced by up to 40% (Aquasar, IBM Zurich) |
| | ▪ Minimize heat exhaust with optional Rear Heat Exchanger |
| **Flexible** | ▪ Excellent flexibility in node and rack configuration |
| | ▪ Factory integrated racks delivered to the data center |

[10] Cabot Partners, "**The IBM System x iDataPlex dx360 M4: Superior Energy Efficiency and Total Cost of Ownership for Petascale Technical Computing", March 2012.**

[11] iDataPlex is available in an air or water-cooled configuration. To put things in perspective, the direct water-cooled iDataPlex system has up to 40% lower power consumption than latest air-cooled iDataPlex. The latest air-cooled iDataPlex system itself has 40% lower power consumption than older air-cooled iDataPlex. (Source IBM)

| | |
|---|---|
| | ▪ 3rd party options and rack support |
| **Affordable** | • Shared infrastructure designs cut cost out |
| | ▪ Non-redundant components |
| | ▪ TCO benefits --- Petascale clusters**: iDataPlex M4 with Intel Xeon processor E5-2600 product family vs. x86 server cluster using Intel Xeon processors 5600 series** - Electricity cost ~90% lower, IT Cap Ex ~60% higher but 57% less overall TCO with iDataPlex as compared to typical x86 based cluster. |
| **Manageable** | ▪ Blade like thinking in design |
| | ▪ Scalable system management software for large-scale configurations |
| | ▪ Rack management appliance for several racks |
| | ▪ Ease in servicing from complete front access for upgrades |
| | ▪ RAS advantages |

**Figure 5: iDataPlex Features and Benefits at a Glance (Source: IBM)**

## Faster results with Intel Xeon processors E5-2600 series

The 32 nm Intel Xeon processor E5-2600 series is the latest processor that replaces platforms based on the Intel Xeon processor 5500 & 5600 series (Nehalem EP and Westmere) that powered IBM's earlier iDataPlex dx360 M2 & M3 servers. The latest IBM's iDataPlex dx360 M4 servers deploy Intel's super-efficient and latest Xeon E5-2600 processors. This brings game changing performance for real world HPC applications and most complex technical problems.

Xeon E5-2600 processors powering dx360 M4 servers provides core count (up to 8 cores), cache size (up to 20MB) and supports more efficient instructions with Intel Advance Vector Extensions (AVX) technology, delivering up to an average of 80% more performance across a range of commercial and technical workloads at consistent power level[12]. Beside the raw performance gains, this the industry's first processor with PCI Express (PCIe) 3.0 I/O integration that allows improved I/O through latency reduction by as much as ~30%, thus reducing network and storage bottlenecks in HPC applications. New fabric controllers that harness PCIe 3.0 specification allow more efficient scaling of performance and data transfer with the growing number of nodes in very large-scale HPC supercomputers.

Improved memory bandwidth and increased cache in E5-2600 processors further boosts the aggregate performance for HPC workloads over prior Intel generations by as much as 70% to 110% depending on the application. In fact, for LINPAC, the E5-2600 processors deliver raw performance of over 151GF/socket, almost 2.2 times its predecessor – the Intel Xeon processor 5600 series. Today, Intel Xeon processor E5-2600 processors power several systems on the TOP500 list of supercomputers.

Besides this, Intel's Quick Path Interconnect (QPI) delivers a substantial increase in bandwidth from a scalable-shared memory by incorporating an integrated DDR3 memory controller onto the processor die. These processors can run two threads per core simultaneously with Intel Hyper-Threading technology, which improves throughput for most HPC workloads. Multi-level shared

[12] Intel press release http://newsroom.intel.com/community/intel_newsroom/blog/2011/11/15/intel-reveals-details-of-next-generation-high-performance-computing-platforms

cache reduces latency to frequently used data thereby improving performance and efficiency significantly. The E5-2600 series processor has integrated memory controller, QPI and integrated I/O.

Intel Turbo Boost Technology 2.0 increases performance of both multi-threaded and single threaded workloads by over 10%[13]. This technology is activated when the Operating System requests the highest processor performance state. The maximum frequency of Intel Turbo Boost Technology is dependent on the number of active cores. While the length of time the processor spends in the Intel Turbo Boost Technology 2.0 state depends on the workload and operating environment, it provides the performance a user needs, when and where the user or application needs it.

For a given workload, the number of active cores, estimated electricity consumption, or processor temperature can set the upper limit of Intel Turbo Boost 2.0 technology. When the processor is operating below these limits and the user's workload demands additional performance, the processor frequency will dynamically increase by 133 MHz on short and regular intervals until the upper limit is met or the maximum possible upside for the number of active cores is reached. Conversely, when any of the limits are reached or exceeded, the processor frequency will automatically decrease by 133 MHz until the processor is again operating within its limits with the stated frequency as a lower bound. For example, a database query, which does not use any of the processor's floating point silicon, can take advantage of that thermal headroom and increase the frequency of all available cores.

### The iDataPlex: Tailored for HPC Applications

HPC users have always demanded computing solutions that have the best performance, price/performance, and now are increasingly demanding energy-efficient platforms. iDataPlex server with its economical, large scale-out architecture is an excellent workhorse for a wide-range of HPC applications in engineering, life sciences, upstream petroleum, financial services, and scientific research. Often, these applications do not require hardware redundancy for failure protection, and software fault-tolerance built into the application and/or system software is more than adequate for practical purposes.

Over the last decade, in conjunction with application developers, IBM has made substantial investments to optimize and tune HPC applications to take advantage of the IBM Cluster architecture. Together they have achieved substantial performance tuning by careful load balancing, maximizing single socket performance, maintaining data locality, minimizing cache misses, and maximizing the computation-to-communication ratio. In many cases, application performance is further enhanced by the use of optimized mathematical libraries and high-performance I/O solutions like the IBM General Parallel File System (GPFS™). All these investments translate readily to iDataPlex dx360 M4 servers.

## 4. iDataPlex Case Studies

IBM's iDataPlex HPC deployments across many industries help reduce the computational costs for Oil & Gas, Financial Services, Space Research, Climate Modeling, Academia and many more. In this section, we present a select set of iDataPlex case studies highlighting the client business challenges, the kind of iDataPlex solution deployed and the benefits gained by the respective IBM clients.

[13] http://www.tomshardware.com/reviews/core-i7-3960x-x79-sandy-bridge-e,3071.html

## NASA Center for Computational Sciences – High Profile Global Climate Research

The NASA Center for Computational Sciences (NCCS) based at the Goddard Space Flight Center in Maryland, is one of two organizations that support the NASA High-End Computing (HEC) Program. The HEC Program provides more than 77,000 computer processors and peak processing power greater than 800 teraflops to the NASA user community. The NCCS is a global leader in climate and weather modeling research, providing specialized support to the science community within NASA and making significant contributions to high-profile global research initiatives.

### Business Challenge: Enhancing high-resolution climate modeling on the Discovery supercomputer cloud

NASA's scientists needed to run higher-resolution simulations to improve their understanding of the changing global climate, boosting their contribution to several major global research initiatives. IBM's iDataPlex enhanced computational capabilities allow NASA climate scientists to run high-resolution simulations that reproduce atmospheric features not previously seen in their models, supporting vital global research.

### Solution

IBM's iDataPlex based cluster solution helps NASA push the research envelope and create ever more detailed climate simulations. Targeting a higher level of simulation, the NCCS has added more than 1,000 iDataPlex dx360 M2 water-cooled servers to its Discover cluster, providing 8,256 new cores of computational power and over 24 TB of high-speed memory.

### Benefits

The NCCS generates vast quantities of data from climate and weather models. The iDataPlex solution brings together sensor data from numerous sources and numerical models to produce sophisticated and highly detailed climate simulations. The new computational capabilities enable NASA climate scientists to create simulations that are more detailed. This helps them in reproducing atmospheric features, which were previously not visible in their models[14].

## Tricon Geophysics – Striking Oil with the IBM iDataPlex

Oil and gas companies are under increasing pressure from their stakeholders to find new reserves faster, and this creates constant demand for more compute power in seismic migration. Tricon Geophysics, Inc. is a processing, seismic imaging and reservoir characterization company for the oil and gas industry.

### Business Challenge: Overcoming data center space and cooling barriers

To improve its customer service, Tricon Geophysics needed to upgrade its supercomputing facilities, but faced strict limitations in the available floor space, cooling capacity and electrical power in its data center.

### Solution

Tricon Geophysics selected an IBM System x iDataPlex solution with a Rear Door Heat Exchanger running Red Hat Enterprise Linux as the basis for its new supercomputer. This flexible and highly

[14] NASA Launches new global research http://www-01.ibm.com/software/success/cssdb.nsf/CS/STRD-7XPKUK?OpenDocument&Site=default&cty=en_us

efficient platform enables Tricon Geophysics to process more data and to improve customer service by more accurately identifying probable oil and gas reserves.

**Benefits**

iDataPlex based cluster solution has quadrupled the speed of processing at Tricon Geophysics, enabling the company to process larger and more complex geo-physical data sets within shorter timeframes. Faster computation enables Tricon Geophysics' customers to find new reserves faster and with a greater degree of certainty, minimizing the need for costly exploratory drilling. It consumes 40% less energy[15] than alternative solutions.

## Leibniz Supercomputing Centre (LRZ) builds iDataPlex based Petascale Supercomputer

The Leibniz Computing Center of the Bavarian Academy of Science is a member of the Gauss Centre for Supercomputing (GCS), the German Federal Computing center, a member of PRACE (Partnership for Advanced Computing in Europe). LRZ houses Europe's Petascale supercomputer – SuperMUC. LRZ caters to diverse interests, providing massive processing power to investigate challenging questions in science and engineering. Researchers trace the evolution of matter in the universe, design aircraft that can carry infrared telescopes, improve the efficiency of and reduce the noise emitted by jet engines, and investigate the characteristics of viruses.

**Business Challenge: Skyrocketing energy consumption**

Power consumption at LRZ has reached very high levels up to 20Mwatt/h. LRZ required a highly energy efficient Petascale supercomputing system, with technology that could bound high power consumption and reduce operational costs without impacting the compute prowess needed by German scientists and researchers.

**Solution**

LRZ deployed the IBM iDataPlex and Linux based Petascale supercomputer – SuperMUC. It uses a new revolutionary form of high temperature cooling developed by IBM for ETH Zürich with the prototype system Aquasar. Active components like processors and memory are directly cooled with water that can have a temperature of up to 45 degrees Celsius. When completed, SuperMUC will consist of more than 9324 nodes with two Intel Xeon E5-2600 series CPUs, 207 nodes with four Intel Xeon processor 5600 series CPUs and over 157,464 cores, ~3PetaFLOP/s peak performance, 324 TB memory and 10-petabyte file space based on IBM GPFS with 200GB/s aggregated I/O bandwidth. The innovative direct water-cooling and energy aware scheduling with IBM xCAT and IBM LoadLeveler will ensure that this high performance cluster has a PUE ~1.1.

**Benefits**

With help from IBM, the LRZ team has designed an energy efficient cluster that will constrain power dissipation and minimize the environmental impact of supercomputing. The new supercomputer will help researchers achieve results rapidly, even for extremely large workloads. The "High Temperature Liquid Cooling" together with very innovative system software promises

[15] Tricon Geophysics accelerates oil and gas discovery with IBM iDataPlex supercomputer ftp://ftp.software.ibm.com/la/documents/imc/la/commons/outcome/pdf/full-stories/oil.pdf

to cut the costs for cooling by 40 percent as compared to air-cooled iDataPlex according to IBM[16]. It will therefore allow for a very moderate increase in energy needed to operate this larger cluster system compared to its predecessor HLRB II. In addition, LRZ will reuse this energy for heating its buildings. Another benefit gained by LRZ is ease of maintenance of iDataPlex cluster with individually serviceable servers, and front access hard drives/cabling. In addition to this, they now have system configurability for customer-specific compute, storage, or I/O needs and pre-configured options for rapid deployment. Besides these, the common tools across the System x portfolio for management at the node, rack or data center level also simplify management of large-scale clusters. With this iDataPlex cluster, LRX will reuse 90% of the waste heat.

## Vestas – Turning Climate into Capital with Big Data

Wind energy is one of today's most important renewable energy sources. By 2020 estimates[17] are that as much as 10 percent of the world's electricity consumption will use wind energy, bringing the wind power industry on par with oil and gas. Precise placement of a wind turbine can affect its performance and its useful life. For Vestas, the world's largest wind energy company, gaining new business depends on responding quickly and delivering business value.

### Business Challenge: Cost efficient expansion of compute capability

Vestas required a cost effective expansion of compute power so that its wind library could scale more than 10-fold to include a larger range of weather data over a longer period, and they needed a more powerful computing platform to run global forecasts much faster.

### Solution

Vestas uses one of the largest supercomputers worldwide along with a new big data modeling solution to slice weeks from data processing times and support 10 times the amount of data for more accurate turbine placement decisions. Improved precision provides Vestas customers with greater business case certainty, quicker results and increased predictability and reliability in wind power generation. The IBM InfoSphere BigInsights software running on an iDataPlex system serves as the core infrastructure to help Vestas manage and analyze weather and location data in ways that were not previously possible. For example, the company can reduce the base resolution of its wind data grids from a 27x27 kilometer area down to a 3x3 kilometer area (about 1.8x1.8 miles)—a nearly 90 percent reduction that gives executives more immediate insight into potential location. Today, the platform that drives its forecasting and analysis consists of a hardware stack based on iDataPlex supercomputer. Such a supercomputing solution enables the company to use 40 percent less energy while increasing computational performance. Each of the system's 12 racks can run twice the number of servers—reducing the amount of floor space required in its data center.

### Benefits

iDataPlex based cluster solution helped Vestas reduce response time for wind forecasting information by approximately 97 percent, from weeks to hours, thus helping them cut development time. It also helped Vestas improve accuracy of turbine placement with capabilities for analyzing greater breadth and depth of data. In addition, it helped to lower the cost for Vestas' customers per

[16] IBM iDataPlex based SuperMUC at LRZ http://www.lrz.de/services/compute/supermuc/systemdescription/

[17] Lars Christian Christensen, VP Vestas Wind System A/S http://www-01.ibm.com/software/success/cssdb.nsf/CS/JHUN-8MY8YK?OpenDocument&Site=gicss67eu&cty=en_us

kilowatt-hour of energy produced and increased their customers' return on investment. Vestas was able to reduce its IT footprint and costs, and decreases energy consumption by 40 percent—all while increasing computational capability with iDataPlex solution[18].

### University of Toronto SciNet cools with iDataPlex

**SciNet** is a consortium of the University of Toronto and affiliated Ontario hospitals. SciNet is one of seven consortia that comprise Compute/Calcul Canada, a national high performance computing resource for academic institutions.

#### Business Challenge

HPC initiatives at the University of Toronto include research in aerospace, astrophysics, bioinformatics, chemical physics, climate change prediction, medical imaging and the global ATLAS project, which is investigating the forces that govern the universe. In order to become world's premier computational research facility, the University of Toronto required a high performance cluster that was energy efficient for a given price and performance.

#### Solution

The University of Toronto's SciNet Supercomputer Datacenter combines massive power, with over 30,000 cores (it's number 67 on the Top 500 list[19]), with exceptional energy efficiency. It has a Power Usage Effectiveness (PUE) rating of 1.16[20] and features an unusual design[21] combining iDataPlex server and the iDataPlex Rear-Door Heat eXchanger, which allows the university to run high-density cabinets with minimal room air conditioning.

#### Benefits

The University of Toronto's SciNet consortium was able to create a powerful supercomputer in Canada to help research in various disciplines. The SciNet machine will be used to create high-resolution global climate models to predict the effects of accelerating melting of Arctic sea ice. One of the first projects that IBM's iDataPlex system will tackle is to create a climate simulation for the province of Ontario and the surrounding Great Lakes region showing the effects of the ice melt. iDataPlex supercomputing cluster at University of Toronto enhanced SciNet's competitive position in globally important research projects.

## 5. Into the Future: Leading HPC with the IBM iDataPlex

High Performance Computing was once the exclusive province of government and premier academic institutions and large enterprises. Today, with technology innovations and evolution, HPC is established well across several industries and in small and medium businesses. Over the last decade, with the widespread penetration of industry-standard clusters, HPC capital expenses as a percentage of IT spend have decreased. However, the associated operational expenses to manage these higher computing density HPC data centers have escalated largely because of increased costs in systems administration, managing RAS, energy, and cooling requirements as systems become denser with thousands of cores and rising facilities costs.

---

[18] Vestas deploys iDataPlex based solution http://www-01.ibm.com/software/success/cssdb.nsf/CS/JHUN-8MY8YK?OpenDocument&Site=default&cty=en_us

[19] Top 500 List, November 2011 http://i.top500.org/site/3049

[20] Compare with industry average PuE of 1.8 according to Uptime Institute http://www.datacenterknowledge.com/archives/2011/05/10/uptime-institute-the-average-pue-is-1-8/ Lower is better.

[21] SciNET cools with iDataPlex http://www-03.ibm.com/press/us/en/pressrelease/27755.wss

The latest Intel Xeon E5-2600 processor powered iDataPlex M4 and the rear door heat exchanger feature along with the latest innovative direct water-cooled iDataPlex will lower data center TCO. These systems will help contain the escalating operational expenses in HPC data centers while reducing capital expenses and protecting a customer's investment in applications, skills, and programming models by as much as 55% as compared to equivalent commodity server based Petascale clusters using Intel Xeon processor 5600 series[22].

With the further mainstreaming of HPC and emerging applications in multimedia, video analytics, voice recognition, biometric security, and other business intelligence and Big Data Analytics domains, HPC spending will grow faster than the overall IT market for servers and storage. The inevitable rising energy costs and increased government regulation that drive greener IT solutions in the industry should ensure a pivotal role in the future for system architectures such as iDataPlex server powered by the latest generation of Intel Xeon E5-2600 processors.

We expect that IBM will continue to evolve newer generations of HPC scale-out clusters by integrating the latest industry-standard, multi-core processors from Intel, and by leveraging IBM's vast experience and IP in large-scale enterprise systems and data centers, particularly in packaging and power and cooling. Some possible future system innovations could be in the design of more effective power supplies, better cooling fans, and more efficient liquid cooling. Possible software innovations for better energy management coupled with systems management, and potential deployment process innovations for next generation data centers as part of IBM HPC cloud computing initiatives could offer HPC customers a unique value proposition of affordability, flexibility, scale, and simplification well into the future.



[22] Cabot Partners, "**The IBM System x iDataPlex dx360 M4: Superior Energy Efficiency and Total Cost of Ownership for Petascale Technical Computing", March 2012.**