The Hartree Centre leads the way in data-intensive computing

Overview

The need

The Hartree Centre has been established in association with IBM as part of a major UK government investment in business-focused high-performance computing. It needed a powerful, energy-efficient infrastructure to support both general computing and Big Data projects.

The solution

The Hartree Centre deployed Blue Wonder, an IBM® System x® iDataPlex® cluster comprising 8,192 Intel Xeon E5-2670 processor cores. Software from IBM and ScaleMP enables shared file system access across the entire cluster and the rapid provisioning of large-scale shared memory environments, ideal for Big Data workloads.

The benefit

Provides an easy entry-point for high-performance computing for UK organisations of all sizes. Supports the development of business applications capable of taking advantage of HPC architectures. Complies with "Green IT" objectives and saves on electricity costs with energy-efficient design and energy-aware scheduling of tasks.

Founded in 2012, the Hartree Centre – a research collaboration in association with IBM – is part of a new centre of expertise in computational science and engineering located at the Daresbury Science and Innovation Campus near Manchester. The centre was created by the Science & Technology Facilities Council (STFC) as a result of a £37.5 million investment by the UK government. Its main objective is to develop a centre for research into business-focused high-performance computing.

Professor John Bancroft, Project Director of the Hartree Centre, explains: "High-performance computing is becoming increasingly cost-effective, and is now within the grasp of even mid-sized businesses. However, very few of the business applications that are on the market today are capable of taking advantage of massively parallel clustered computing architectures. The main objective of the Hartree Centre is to give the UK a head-start on the development of new applications that will be able to exploit HPC infrastructures, and to give businesses of all sizes an easy entry-point to the world of HPC."

Multi-purpose solution

The Hartree Centre is building several HPC clusters to support various different types of research project. One of the most important requirements was to create a general-purpose cluster built on familiar, industry-standard technologies such as x86 processors and Linux. This would provide an easy introduction to HPC technologies for businesses that have no prior experience of the challenges and opportunities of cluster computing. A second requirement was to create a large shared-memory cluster that would support research into Big Data challenges such as analysing weather and climate data or performing complex aerodynamic modelling for the aviation and automotive industries.

"We realised that it might be possible to meet both of these requirements with a single multi-purpose cluster," comments Professor Bancroft. "We received bids from IBM, Fujitsu, SGI and T Platforms for the hardware, and various processor architectures were proposed including Intel, AMD and NVIDIA. We chose IBM based on the whole package they offered – the technical architecture they proposed was

Solution Components

Hardware

- IBM® System x® iDataPlex® dx360 M4 with Intel Xeon processor E5-2670
- IBM System Storage® TS3500 tape library with TS1140 Enterprise tape drives
- DataDirect Networks SFA 10K-X disk systems

Software

- IBM General Parallel File System (GPFS™)
- IBM Platform™ HPC
- ScaleMP vSMP software
- Red Hat Enterprise Linux

Services

- IBM Systems and Technology Group Services for HPC

IBM Business Partner

- OCF

as good as anything on the market, and their collaborative approach to the relationship was a perfect fit with our needs. The support they have provided, both directly and through their partner OCF, has been second to none."

OCF is one of the UK's leading specialists in high-performance computing, and has worked with customers in many fields to build supercomputing infrastructures for all kinds of applications. The combination of experience and expertise brought to the table by OCF was a significant factor in the project's success.

Sharing the benefits

The general-purpose and Big Data cluster at the Hartree Centre has been named Blue Wonder – a homage to both the scientific breakthroughs it will help to achieve, and to Lewis Carroll, the author of Alice in Wonderland, who was born in Daresbury.

As a platform for Blue Wonder, the Hartree Centre selected an IBM System x iDataPlex dx360 M4 architecture. The cluster is composed of 512 computing nodes, each of which has a minimum of 32 GB of memory and dual eight-core Intel Xeon processor E5-2670. The entire Blue Wonder system is managed using IBM Platform™ HPC, which provides all necessary functions to effectively deploy, manage and support a complex high-performance computing environment. From within the Platform HPC software, the Hartree Centre can use vSMP software from ScaleMP to dynamically pool processing resources and memory from selected nodes into a single virtual resource pool. With the vSMP software running on 384 of the iDataPlex nodes, this means that a single application could theoretically harness up to 6,144 cores and 48 TB of memory – enabling enormous databases to be loaded into memory and subjected to sophisticated analysis.

If several smaller jobs need to utilise the shared memory architecture, the new cluster can easily accommodate them. IBM Platform HPC makes it easy to dynamically create shared memory domains of different sizes and then return these resources to the shared pool when the jobs finish.

"The flexibility of Blue Wonder is a real step forward from traditional HPC clusters, which are generally designed to be managed by computer scientists, and are not particularly user-friendly," comments Professor Bancroft. "For businesses to embrace HPC technology, it needs to be relatively easy to use. The combination of the iDataPlex and Platform HPC software makes this a reality – you can simply request, for example, 100 cores and 1 TB of memory, and the software provisions the environment automatically. And when you've finished with it, the resources are returned to the pool."

"iDataPlex gives us an ideal platform to help businesses of all sizes to harness high-performance computing and generate real economic impact. We're delighted that IBM, Intel and OCF have shown themselves to be such worthy partners in this endeavour."

— Professor John Bancroft, Project Director at the Hartree Centre

Familiar technology

Another factor in the choice of the iDataPlex platform was its use of the Intel Xeon processor E5-2670 – a versatile processor designed to deliver the best combination of performance, built-in capabilities and cost-effectiveness for a wide range of needs. In addition to being perfectly suited for an HPC cluster that will be used to process multiple different types of workload, the Intel Xeon processor family is ubiquitous in business IT, so it will be a familiar technology for companies who are interested in utilising the cluster.

"If we're developing new software for businesses to use, it's sensible to base it on technologies that most businesses use and understand," comments Professor Bancroft. "Intel have also been very supportive, and have worked closely with IBM and OCF to deploy, configure and test the cluster successfully."

Andy Grant, HPC Sales and Business Development Lead at IBM, adds: "We did some testing internally to gauge the performance of the Intel processors against an alternative from AMD, and in this configuration they achieved better results across a range of representative benchmarks. So we had no hesitation in recommending them to the Hartree Centre."

In fact, tests show that the Blue Wonder iDataPlex cluster can achieve 206.3 teraFLOPS, which brings it in at number 114 in the June 2012 Top 500 list of supercomputers. Its 48 TB shared memory capacity also makes it the largest shared memory cluster in the UK.

"The iDataPlex gives us one of the most powerful clusters of its type in the country, but it's also very energy efficient," says Professor Bancroft. "With its Rear Door Heat eXchanger water-cooling technology and energy-aware scheduling, it scored very well on the electricity consumption portion of our selection criteria. Environmental sustainability is naturally a key priority for any government-funded IT system, and on a more practical level, it saves us money. Spending less on electricity means we can spend more on running projects and funding research – so everybody wins."

Broad applications

Working in conjunction with Daresbury's Virtual Engineering Centre, which focuses on computer-based visualisation, simulation and modelling projects, the Hartree Centre is now working to attract businesses to take advantage of the capabilities of the new cluster.

"The VEC works with many key players in engineering industries – from big names such as BAE Systems, Bentley and Jaguar Land Rover to small and medium niche businesses," states Professor Bancroft. "With its ability to handle massive data-sets in memory, our cluster offers a huge opportunity for advanced manufacturing companies to study highly complex engineering problems – so we're very keen to link up with some of the VEC's partners."

Other likely applications for Blue Wonder include life sciences projects studying molecular modelling, genomic databases and cosmetics design; energy and environment projects studying weather and climate data and nuclear energy; and digital economy projects in a wide range of areas, such as gaming and the analysis of CCTV footage.

Professor Bancroft concludes: "iDataPlex gives us an ideal platform to help businesses of all sizes to harness high-performance computing and generate real economic impact. We're delighted that IBM, Intel and OCF have shown themselves to be such worthy partners in this endeavour."

About OCF

OCF has been using its integration knowledge, skills and partner eco-system to help customers with their high-performance data processing, data management and data storage challenges for over a decade. Based in Sheffield UK, with a network of expert staff around the country, OCF provides solutions to a growing number of commercial clients from the automotive, aerospace, financial, manufacturing, media, oil & gas and pharmaceutical industries. It also provides solutions to over 20 percent of the UK's universities, higher education institutions and research councils.

To learn more about products, services and solutions from OCF, please visit **www.ocf.co.uk**

For more information

To learn more about IBM iDataPlex high-performance computing solutions, contact your IBM sales representative or visit **ibm.com**/systems/info/x/idataplex

To learn more about IBM Platform Computing, contact your IBM marketing representative or IBM Business Partner, or visit **ibm.com**/platformcomputing

IBM United Kingdom Limited
PO Box 41
North Harbour
Portsmouth
Hampshire
PO6 3AU

Produced in the United Kingdom
August 2012

XSC03129-GBEN-01