HPC Profiles in Leadership

The University of Birmingham (UK)

Steve Conway and Kevin Monroe
June 2017

HYPERION RESERCH OPINION

The University of Birmingham in the city's Edgbaston area is the United Kingdom's fourth-largest university, serving about 34,000 undergraduate and graduate students, along with researchers from throughout the UK and beyond.

The university's Computer Centre, to the north of the Edgbaston campus, is the centerpiece of the Birmingham Environment for Academic Research (acronym: BEAR), a collection of IT resources that are available without cost to the University of Birmingham community and qualified external researchers. At the heart of the BEAR environment is the BlueBEAR high performance computing (HPC) cluster.

The upgrade to BlueBEAR3, now in progress, debuts cooling technology that was co-developed by system vendor Lenovo and the staff of the university's Research Computing Team.

- This innovative technology is expected to cut cooling-energy costs by up to 83 percent and permit substantially denser server configurations—benefits that Lenovo plans to exploit at other customer installations, including the University of Birmingham's own BEAR Cloud HPC distributed research environment and Lenovo's global HPC innovation center in Stuttgart.

Hyperion Research applauds the University of Birmingham and Lenovo for extending the prior university-IBM R&D collaboration to produce an energy-saving innovation that promises to benefit both collaborators and the larger HPC community. As an added bonus, the university will deploy the same server and cooling technology in its HPC data center and HPC cloud environment to provide a familiar, seamless resource to researchers across many disciplines.

Note: this page is intentionally blank.

IN THIS HPC PROFILES IN LEADERSHIP REPORT

This profile looks at an innovative new HPC system being installed by Lenovo at the University of Birmingham. The innovation greatly reduces cooling costs and provides for higher density installations.

SITUATION OVERVIEW

The University of Birmingham in the city's Edgbaston area is the United Kingdom's fourth-largest university, serving about 34,000 undergraduate and graduate students, along with researchers from throughout the UK and beyond. The university's Computer Centre, in the Edgbaston Green Zone, is the centerpiece of the Birmingham Environment for Academic Research (acronym: BEAR), a collection of IT resources that are available without cost to the University of Birmingham community and qualified external researchers. At the heart of the BEAR environment is the BlueBEAR high performance computing (HPC) cluster.

The upgrade to BlueBEAR3, now in progress, debuts cooling technology that was co-developed by system vendor Lenovo and the staff of the university's Research Computing team. This innovative technology is expected to cut cooling-energy costs by up to 83 percent and permit substantially denser server configurations—benefits that Lenovo plans to exploit at other customer installations, including the University of Birmingham's own BEAR Cloud HPC distributed research environment and Lenovo's global HPC innovation center in Stuttgart. Hyperion Research applauds the University of Birmingham and Lenovo for extending the prior university-IBM R&D collaboration to produce an energy-saving innovation that promises to benefit both collaborators and the larger HPC community.

As an added bonus, the university will deploy the same server and cooling technology in its HPC data center and HPC cloud environment to provide a familiar, seamless resource to researchers across many disciplines.

WHY HPC IS IMPORTANT TO THE UNIVERSITY

Serving A Diverse User Community

The university's HPC user community is highly diverse. It includes researchers from physics, chemistry and engineering to medicine, economics, archeology, theology and other domains. A few examples:

- During the past two years, many new life sciences researchers have joined the university, creating a large increase in demand for HPC and a need for domain-specific software stacks. Projects range from bioinformatics for public health to research on virus evolution and the Zika virus, epidemiology, and translational genomics.
- Through research partnerships, commercial organizations perform substantial simulation on the university's supercomputer resources.
- Linguistics researchers use the supercomputer resources to study how languages change over time, borrowing bio-classification methods in their work.

- Psychology and robotics researchers collaborate to advance the state of the art in machine learning, deep learning and robotics.
- Other researchers exploit the supercomputer for archeology studies.

BLUEBEAR3: THE UPGRADE GOES MODULAR

Today's 29.3TF BlueBEAR2 Linux system is based on IBM's iDataPlex design incorporating Intel Sandy Bridge ES-2660 dual-processor, 8-core nodes. The BlueBEAR3 upgrade, reflecting the 2014 sale of IBM's x86 server business to Lenovo, is based on Lenovo NeXtScale dense servers and features these specifications:

- 34 nodes, each with two 12-core ES-2690v3 Haswell sockets running at 2.6GHz and 128GB RAM, with a potential boost to 3.5GHz. The total for this configuration is 782 cores (24 cores per node times 34 nodes), with an allocation of 5.3GB RAM per core.
- 58 additional nodes, each with two 10-core ES-2640v4 Broadwell sockets running at 2.4GHz, with a potential boost to 3.4GHz. The total here is 1160 cores (20 cores/node times 58 nodes), with an allocation of 2.8GB RAM per node.
- 2 air-cooled nodes, each with two 10-core ES-2640v4 Broadwell sockets running at 2.4GHz, with a potential boost to 3.4GHz, are equipped with NVIDIA P100 GPGPU cards to provide a base level of access to GPGPU compute resources.
- For more memory-intensive jobs, BlueBEAR3 also includes smaller numbers of ES-2640v4 Broadwell nodes with 256GB RAM and 512GB RAM, respectively.
- All nodes are connected via 100Gb/second EDR InfiniBand. The operating system is CentOS 7.3. Storage is IBM Spectrum Scale provided by Lenovo, running on their DSS-G platform.

As at many other major research universities throughout the world, there is healthy tension at the University of Birmingham between centralized and distributed IT resources, including HPC systems. Research grants can include budgets to acquire clusters for dedicated use by the research teams—a factor that normally weighs against turning to the centralized HPC data center. But the university's modular design approach means that the Computer Centre can incorporate departmental clusters and manage them efficiently as resources reserved for their users alone.

THE BEAR CLOUD ENVIRONMENT

The university describes BEAR Cloud as a collection of "powerful computing resources (dedicated to supporting computationally or data intensive research) that can be deployed flexibly and changed easily to meet evolving user needs. The entire infrastructure is housed on campus; a characteristic which is particularly important for those dealing with large datasets or sensitive data. It is also highly cost effective in comparison with current commercial cloud offerings."

BEAR Cloud was launched in October 2016 with a major investment from the university, including £2 million for life sciences. (The resources exclusively reserved for research in the life sciences are known as CaStLeS.) The private cloud environment, like BlueBEAR3, is based on Lenovo NeXtScale servers. Unlike its BlueBEAR on premise counterpart, however, BEAR Cloud offers the ability to create multiple virtual servers on demand; all co-existing on the same infrastructure. These servers can have different operating systems, such as different Linux variants, and can be tailored for applications that have been developed for specific environments.

The university says "there are currently no commercial cloud offerings that can compete with BEAR Cloud in terms of performance and integration with campus services." In addition, university researchers

carry out extensive medical research and don't feel comfortable yet entrusting this to public cloud security.

The University specifically selected the Mellanox ConnectX-4 card for fabric connectivity as it enables both Infiniband and Ethernet capability in a single adapter.

Operations and Support

The BlueBEAR3 HPC supercomputer and the BEAR HPC cloud are operated by the university's Research Computing team, with first-level support provided by systems integrator OCF and second-level support from Lenovo. Sheffield-based OCF (founded 2002) is an important Lenovo HPC partner in the UK, having gained years of experience in the same role for IBM. OCF teams with the OEM system vendor at other major sites, including the University of Bristol, the University of Oxford, the Edinburgh Parallel Computing Centre (EPCC) and the Hartree Centre in Daresbury. OCF has taken UK customers to Lenovo's global HPC Innovation Center in Stuttgart.

INNOVATIONS: THE NEXTSCALE COOLING SYSTEM

Both BlueBEAR 3 and the BEAR Cloud environment will feature the Lenovo NeXtScale direct on-chip warm water cooling technology. The university and Lenovo co-developed the environment leveraging this innovative cooling technology that Lenovo believes could cut the university's cooling-energy consumption by up to 83% compared with air cooling alone.

The NeXtScale technology flows water at up to 45°C. (113°F.) into the rear of the server via heat sinks attached to the CPUs, dual in-line memory modules, I/O and other components. Water returning from the components withdraws heat from the system, rising in temperature by about 10°C. (18°F.) in the process. The water contains anti-corrosion and anti-bio agents and operates in a closed-loop circulation, with a heat exchanger and pump located outside of the data center building.

Lenovo will exploit the cooling technology in the company's global HPC innovation center in Stuttgart, Germany, and in other deployments of its NeXtScale dense server systems.

Benefits: Dense Servers with Energy-Efficient Cooling

Superior cooling ability makes it possible to operate dense server systems, i.e., servers with greater heat dissipation needs. That means the BlueBEAR3 supercomputer with NeXtScale dense servers and warm water cooling brings these important benefits:

- Substantially increased computing capacity within the Computer Centre's existing spatial and power envelopes.
- Substantially increased performance, not only because of the increased computing capacity, but also because greater density enables faster data movement–data typically has shorter distances to traverse.
- Substantially reduced power costs, as noted above.

It will be interesting to see how well these benefits are exploited on behalf of users' real-world applications. In view of the University of Birmingham's Research Computing Team's positive track record, Hyperion Research is confident that the university's user community will find BlueBEAR3 and its BEAR Cloud counterpart very satisfactory resources.

OPPORTUNITIES AND CHALLENGES

Opportunities

BlueBEAR3's modular design creates multiple opportunities:

- The supercomputer can be added to in small, variable installments. This smooths out the capital expenditure budget, obviating the need for a major capex bump every 3-4 years to upgrade the whole system at once.
- Hardware and software configurations of special interest to certain research teams can be added more readily, with less concern about their impact on the design and limited budget for a whole-system upgrade. Promising new technologies can also be more easily acquired in small doses, to promote early experimentation.
- As at many other major research universities throughout the world, there is healthy tension at the University of Birmingham between centralized and distributed IT resources, including HPC systems. Research grants can include budgets to acquire clusters for dedicated use by the research teams—a factor that normally weighs against turning to the centralized HPC data center. But the university's modular design approach means that the Computer Centre can incorporate departmental clusters and manage them efficiently as resources reserved for their users alone.

Challenges

The modular design also creates some challenges:

- A notable challenge is that modular, incremental upgrades produce HPC systems with multiple generations of CPUs and other hardware/software components and technologies, as has happened with BlueBEAR. This can make programming and system management more difficult.

ESSENTIAL GUIDANCE FOR OTHER HPC SITES

- **Weigh the benefits of liquid cooling**. Recent Hyperion Research studies confirm that the combination of dense servers and liquid cooling (whether with water or another suitable fluid) is increasingly popular as HPC sites attempt to squeeze more computing capacity into the limited spatial confines of their data centers. As the University of Birmingham examples illustrates, HPC data centers that were not constructed with high raised floors and other requirements for traditional liquid cooling deployments have other options for taking advantage of liquid cooling.
- **Consider modular HPC server system designs.** The main benefits and challenges associated with modular designs are summarized in the preceding section. Hyperion believes that for many HPC sites, the opportunities/benefits will outweigh the challenges—especially the ability to spread capex budgeting more evenly across fiscal years.

Adopt the same baseline technologies for on premise and private cloud resources. The University of Birmingham's decision to use the same baseline server and storage systems for the on-premise BlueBEAR3 supercomputer and the BEAR Cloud environment promises to simplify management in many dimensions—routine maintenance, patching, spares, additions, programming and others—while still allowing these two resources to employ very different configurations, software variants, policies and directive to accommodate the varying requirements of their users. We encourage other HPC sites to evaluate whether this strategy makes sense for them.

About Hyperion Research, LLC

Hyperion Research, consisting of the former Hyperion high performance computing (HPC) analyst team, provides HPC information, analysis, and recommendations based on technology and market trends. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.hpcuserforum.com
