

Platform LSF

Helping Beat Cancer with HPC



Customer

Cancer Research UK

Industry

Life Sciences

Challenges

- Having installed a new genetic sequencer to advance its study of common cancers, Cancer Research UK's Cambridge Research Institute (CRI) needed to manage the five terabytes of data it produces each week in order to carry out new research into tumour progression
- As a charitable organization, it was important for CRI to deploy a solution that is cost effective so that they can channel their funds to their most important resource: their research scientists.

Solution

Platform LSF – The HPC Workload Management Standard

Results

- Higher cluster efficiency means that Platform LSF's cluster management is opening new fields of research for CRI, helping it in its work to understand the progressive changes of cancers and their resistance to treatments
- Minimal support or manual intervention is required to keep the workload management system operating, reducing staff requirements and allowing CRI to save man hours to the equivalent of one full time employee during the first year
- The computing cluster management capability handles application workload intelligently by balancing demand with available computing resources, so that the vast amounts of research data being produced does not overwhelm CRI's IT systems
- Platform LSF is successfully scaling up as CRI's workload grows, enabling it to future-proof its research requirements
- Cluster solution delivers around .£50,000 in cost savings by removing the duplication of resources across research groups

CRI, a world leader in the study of cancer, is home to over 450 researchers and staff. The organisation specialises in practical high-quality research addressing issues surrounding cancer diagnosis, treatment and prevention. In order to stay at the forefront of global cancer research it is vital for the organisation to be able to expand into new areas of study, regardless of the amount of raw data that this might generate.

CRI recently installed a next-generation sequencing platform with the aim of studying the genetic changes that manifest as a cancer evolves. Understanding these genetic changes are key to both evaluating how a cancer progresses and how it reacts to various medical treatments. The new sequencer, however, was simply producing too much data for CRI's IT systems to process. It was clear that for the data produced by the sequencer to be of any value CRI needed to upgrade its infrastructure with greater processing capacity.

"Platform LSF has enabled us to undertake new research studies that otherwise would not have been possible."

Peter Maccallum
Head of IT & Scientific Computing, Cancer Research UK

A Balancing Act

CRI examined a number of alternative methods of processing the data created by its sequencer. Supercomputers were quickly ruled out on the grounds of cost. As Peter Maccallum, Head of IT & Scientific Computing at Cancer Research UK commented: "keeping costs down is a priority for any IT team, but even more so for CRI. The more money we can save on the cost of our IT, the more we have to spend on our core research projects."

Having also evaluated Open Source solutions, CRI decided to deploy Platform Computing's High Performance Computing (HPC) software. According to Maccallum: "Open Source would have proved too complex and time consuming to install. With Platform LSF we knew we were receiving a solution with a good and proven pedigree that could be implemented instantly and work from the outset. We also knew from our evaluation process that Platform LSF would be easy for us to administer – this was a big factor in us choosing to run with it."

Shared benefits for all

Platform's LSF is a workload management software solution that runs the researcher's applications in an orchestrated manner. Its intelligent workload scheduler utilises a cluster of servers within a single IT data centre location so that all users can share the same servers within CRI's 21 research groups. Previously, each research group has its own computing resources, which drove up server costs, reduced system utilization rates, increased maintenance overheads and complexity. Now with a shared server approach the organisation experiences cost savings of around £50,000 just from removing the duplication of hardware and maintenance across its sites, while also increasing the amount of data being processed.

Since CRI's IT resources are dynamically assigned in-line with the institute's research priorities, it is now able to operate on a leaner IT team for maintaining data processing functions arising from the sequencer. This means the institute now has more available resources to be re-directed to where they are needed most – the research teams. In fact, due to Platform LSF CRI has saved the equivalent in man hours to one full time employee.

The Platform LSF solution also helps ensure that the organisation is meeting its service level agreements (SLA) by guaranteeing resources to research groups or jobs. Its Fairshare Policies allow each group or user to receive their "fair share" of the resources over a defined period ensuring they have the computing power required to conduct their research.

For Maccallum, however, the greatest benefit of Platform LSF has been the new areas of research the HPC solution has opened up. As Maccallum states: "Platform LSF has enabled us to undertake new research studies that otherwise would not have been possible. We now have the means to produce and manage a wealth of gene sequencing data that we could only have dreamed about previously. This has already led to tangible published work looking into breast cancer and is proving its worth in helping our researchers further the understanding of how cancers progress."

Coping with future demand

As a research hub, CRI collaborates with many other research organisations. To enable collaboration on large data sets with high performance clusters, CRI is working with Platform Computing to architecturally support cross-organisation systems. For the time being, however, the institute plans to scale up Platform LSF internally by adding more servers as the computing requirements increase for its valuable work in the fight against cancer.



Platform Computing is the leader in cluster, grid and cloud management software - serving more than 2,000 of the world's most demanding organizations for over 17 years. Our workload and resource management solutions deliver IT responsiveness and lower costs for enterprise and HPC applications. Platform has strategic relationships with Cray, Dell™, HP, IBM®, Intel®, Microsoft®, Red Hat®, and SAS®. Visit www.platform.com.

World Headquarters

Platform Computing Inc.
3760 14th Avenue
Markham, Ontario
Canada L3R 3T7
Tel: +1 905 948 8448
Fax: +1 905 948 9975
Toll-free Tel: 1 877 528 3676
info@platform.com

Sales - Headquarters

Toll-free Tel: 1 877 710 4477
Tel: +1 905 948 8448

North America

New York: +1 646 290 5070
San Jose: +1 408 392 4900

Europe

Bramley: +44 (0) 1256 883756
London: +44 (0) 20 3206 1470
Paris: +33 (0) 1 41 10 09 20
Düsseldorf: +49 2102 61039 0
info-europe@platform.com

Asia-Pacific

Beijing: +86 10 82276000
Xi'an: +86 029 87607400
asia@platform.com
Tokyo: +81(0)3 6302 2901
info-japan@platform.com
Singapore: +65 6307 6590
wliaw@platform.com