DataCentred

SUCCESS STORY

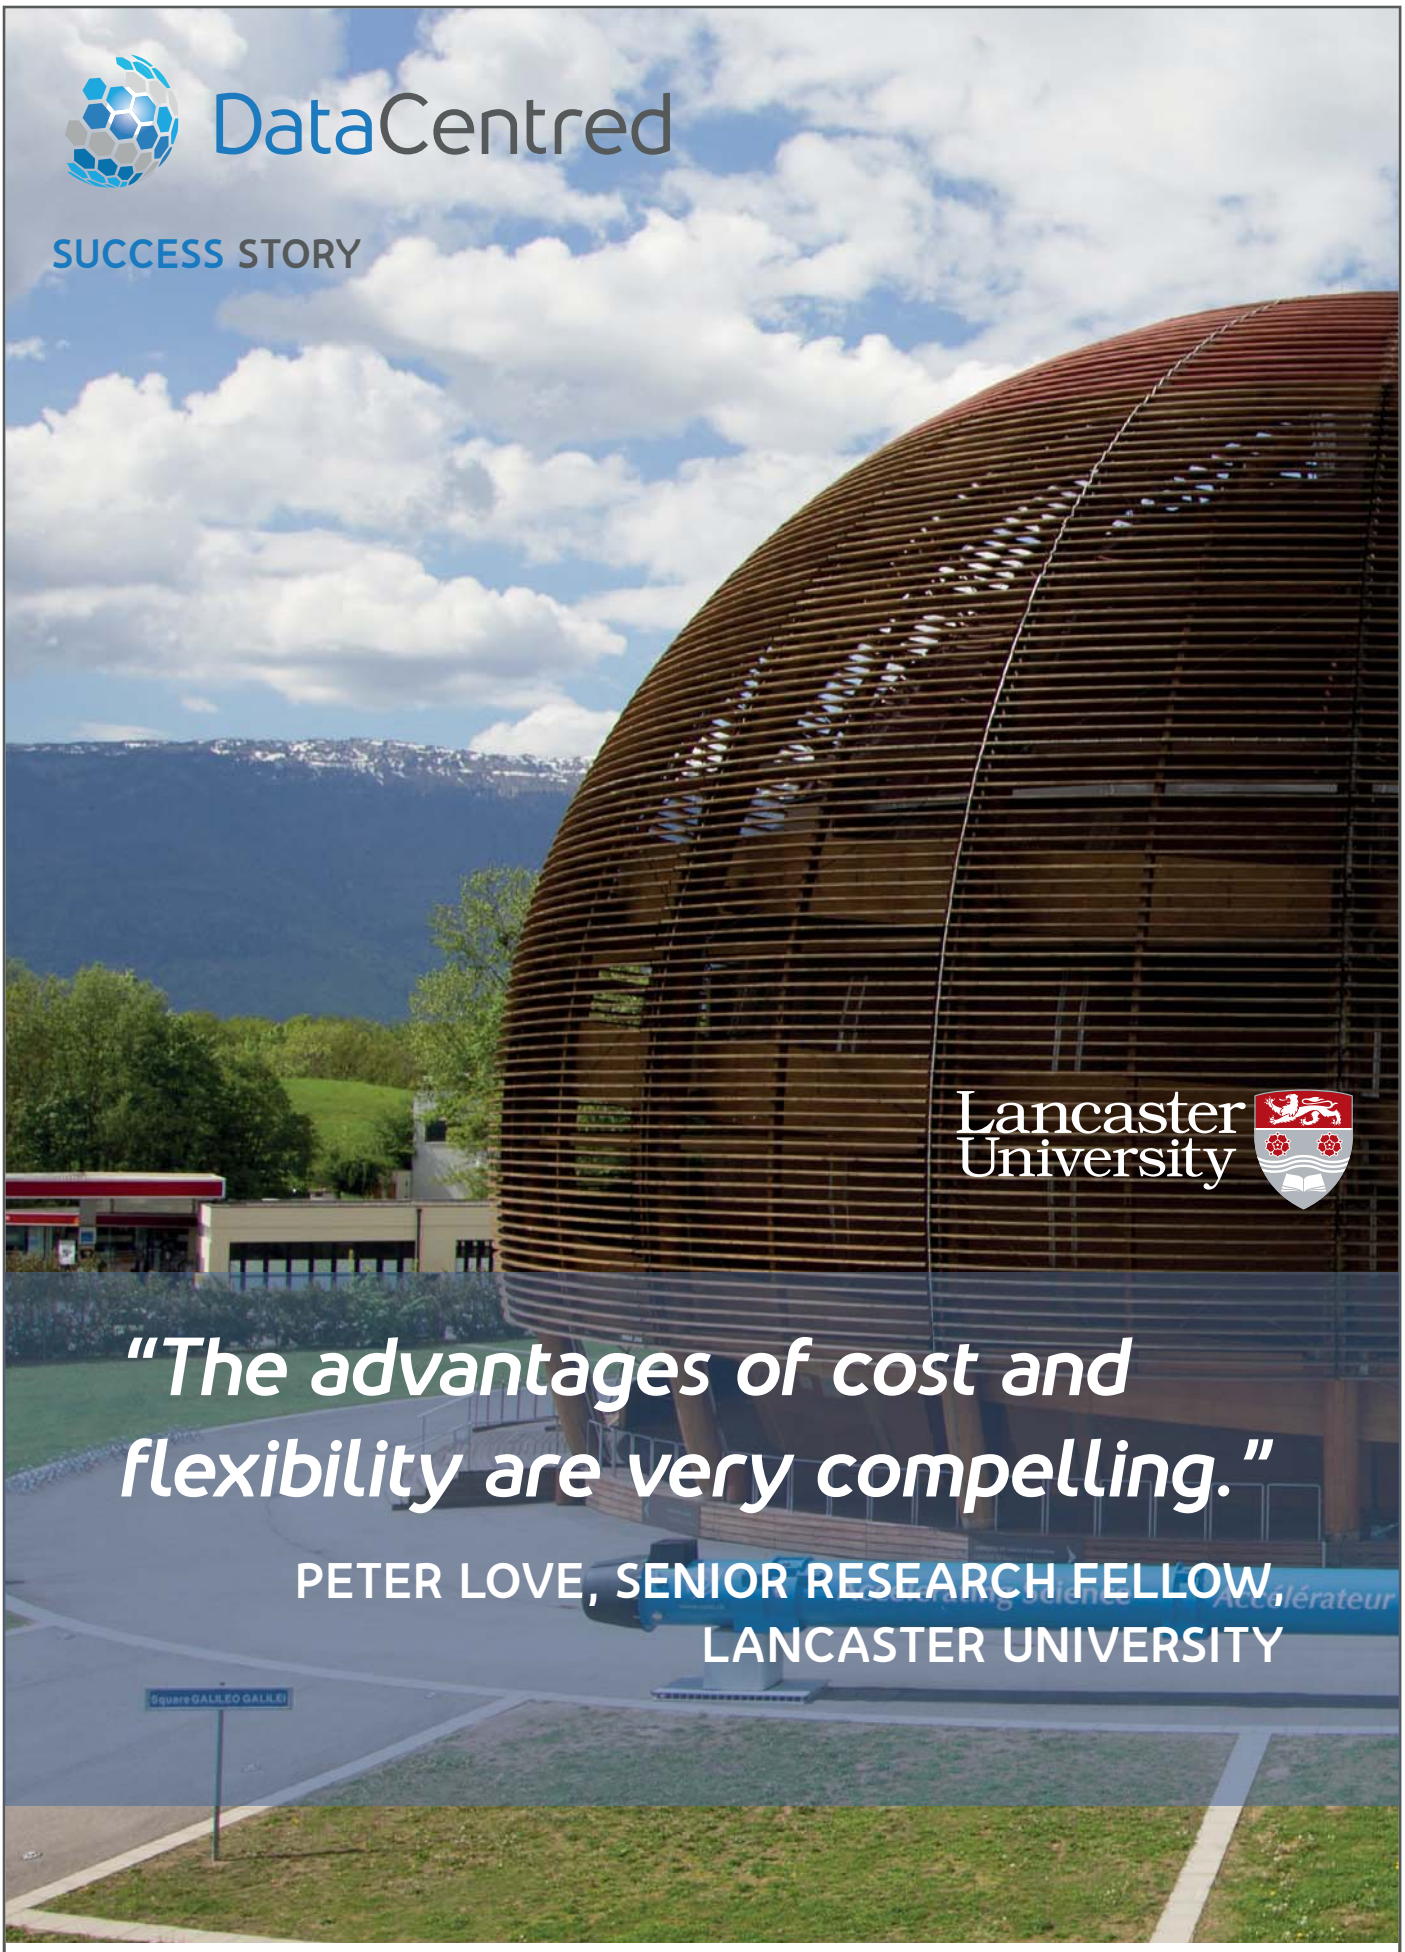

"The advantages of cost and flexibility are very compelling."

PETER LOVE, SENIOR RESEARCH FELLOW,
LANCASTER UNIVERSITY

openstack POWERED · ceph · LINUX FOUNDATION SILVER MEMBER · ubuntu · ARM · HM Government G-Cloud Supplier

DataCentred

Research & Collaboration OpenStack Cloud Compute

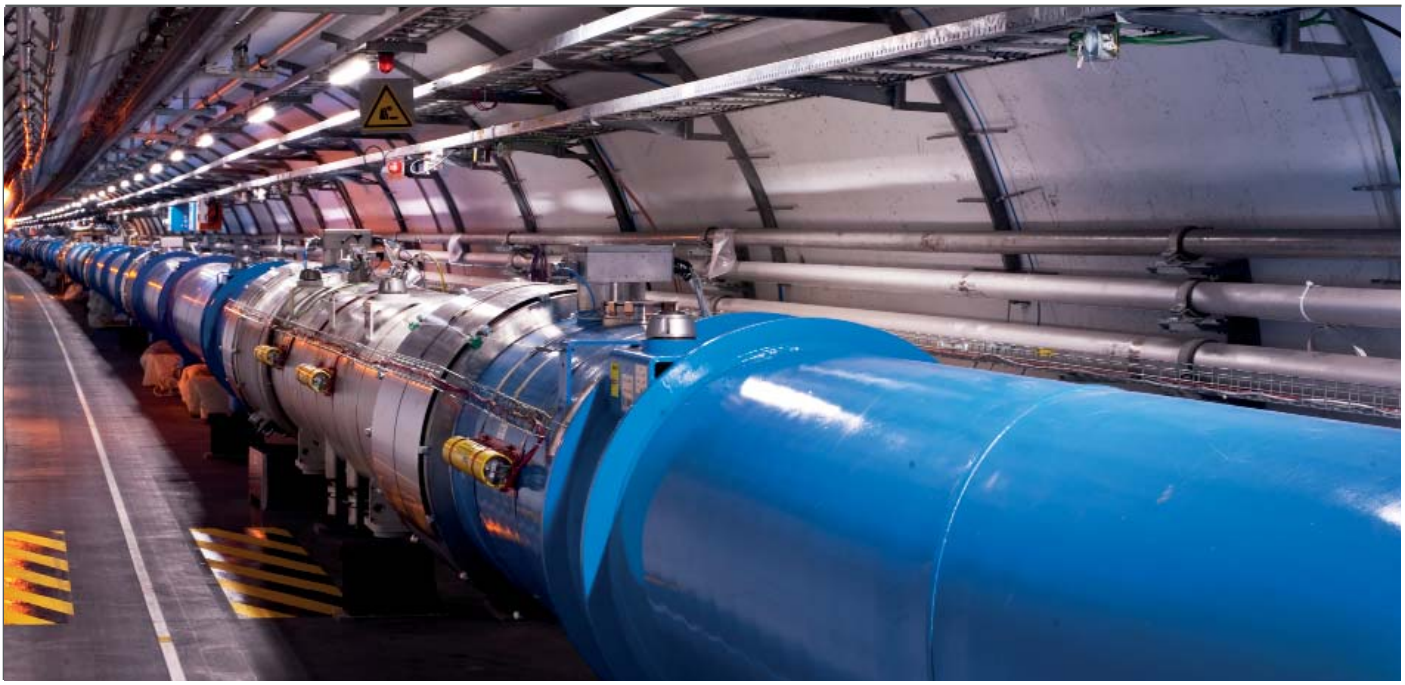

CERN has one of the biggest deployments of OpenStack in the world. So when Lancaster University's Experimental Particle Physics Group sought to move some of its ATLAS workload into the cloud, DataCentred's OpenStack cloud compute platform was a natural first choice.

ATLAS is the largest experiment being conducted at the Large Hadron Collider (LHC) in CERN. Raw data is sent from CERN to regional Tier 1 data centres for storage and processing. From there, it is distributed to the associated Tier 2 centres for further processing. The Experimental Particle Physics Group at Lancaster University operates one such Tier 2 site. It also has responsibility for developing the computing infrastructure – systems and software for processing data – for the ATLAS experiment. This is no small task: thus far, the ATLAS project has stored approximately 150 petabytes of data, both real and simulated. At any one time, the global production system has 150,000 jobs running.

Peter Love is a Senior Research Fellow at the Experimental Particle Physics Group based at Lancaster University. The team at Lancaster were instrumental in establishing the worldwide computing grid for work on the LHC experiments across various university research group sites. Here in the UK, the group started looking into what cloud computing could offer its operations as early as 2012. Peter says that the drivers for the Group investigating the cloud model were threefold:

- The growing role of cloud in the wider IT industry,
- The flexible nature of the cloud,
- The greater elasticity and subsequent cost savings.

Peter explains: "Cloud looked attractive and we wanted to investigate whether it would deliver the promised benefits. A key motivation was the desire to have more flexibility and control of the software stacks deployed at the various sites. Our existing computing infrastructure is a grid system and this doesn't give us much control over the local machines. We have an agreement that it needs to be compatible with our software, but that's about it. The great advantage about using DataCentred's cloud compute is that users can run their own virtual machine and they aren't reliant on the local systems, what operating systems or software has been installed locally."

As well as providing greater flexibility, the elasticity that DataCentred's cloud solutions offered was another key attraction for Peter. The data processing requirements fluctuate every day, as Peter explains: "Cloud computing claims to be elastic enough to allow for fluctuations

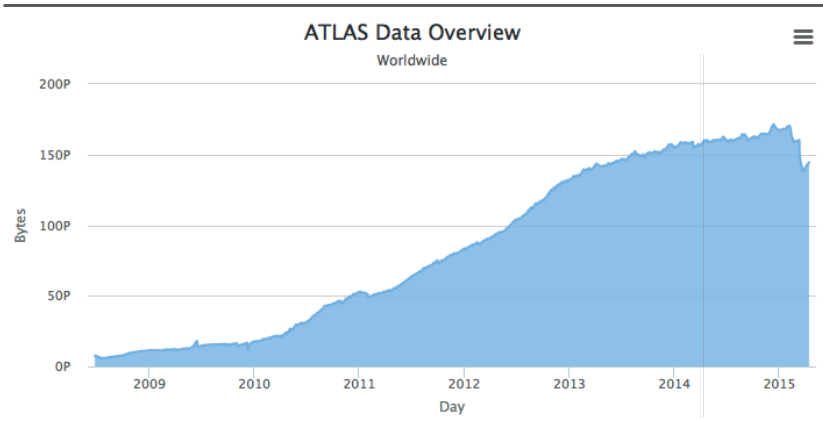

in workload. This is important to us because we do have marked fluctuations. For instance, prior to the big conferences, the physicists will want to run a lot of analyses and the workload will go up. We need to enable this analysis but equally we don't want to be paying for electricity to operate machines when they're not being used. It's wasteful. A major advantage of the cloud model is that you only pay for what you're using."

Having established that cloud computing was worth investigating, the UK ATLAS Computing Group began to investigate the alternatives. Peter says: "We knew that as an open-source software, OpenStack offered real cost advantages. CERN's existing commitment to the software meant that an OpenStack solution would be a lot easier to integrate with our production tools. At that time, there were very few OpenStack cloud providers who were part of the Helix Nebula group. As a member of this group, and a UK-based OpenStack provider, DataCentred was definitely worth investigating. It's wonderful to see local North-West companies engaging with the local scientific community."

The decision to engage with DataCentred was largely based on the team's desire to work with people who are willing to collaborate. Peter says: "Our main criteria for selecting a a commercial partner was a willingness to collaborate and to talk with us. We'd done some development work ourselves and we needed to find someone running an appropriate system, namely OpenStack. It was still early days for us and there had to be a certain dialogue with the engineers on site. We need open back-and-forth communication to get things working. With DataCentred we have genuine communication. From an engineering point of view, they are extremely technically knowledgeable and they are very professional."

"The technical expertise of the DataCentred team has enabled us to share knowledge with other UK institutions."

PETER LOVE

The good communication established between the two sites was especially beneficial early when Peter and his team first started moving the workload onto OpenStack. He says: "At first, there was a lot of learning for us. We'd approach the DataCentred team with our understanding and they would correct us! We had a very quick turnaround of problems. Because of DataCentred's technical expertise they could diagnose a problem very quickly. Instead of getting bogged down in bug hunting, the engineers got to the root cause of the issue and this means a better engineering solution."

This has benefitted the UK ATLAS Computing Group in ways that reach beyond the Lancaster site: "Our research partners at Oxford and Imperial are deploying OpenStack

DataCentred

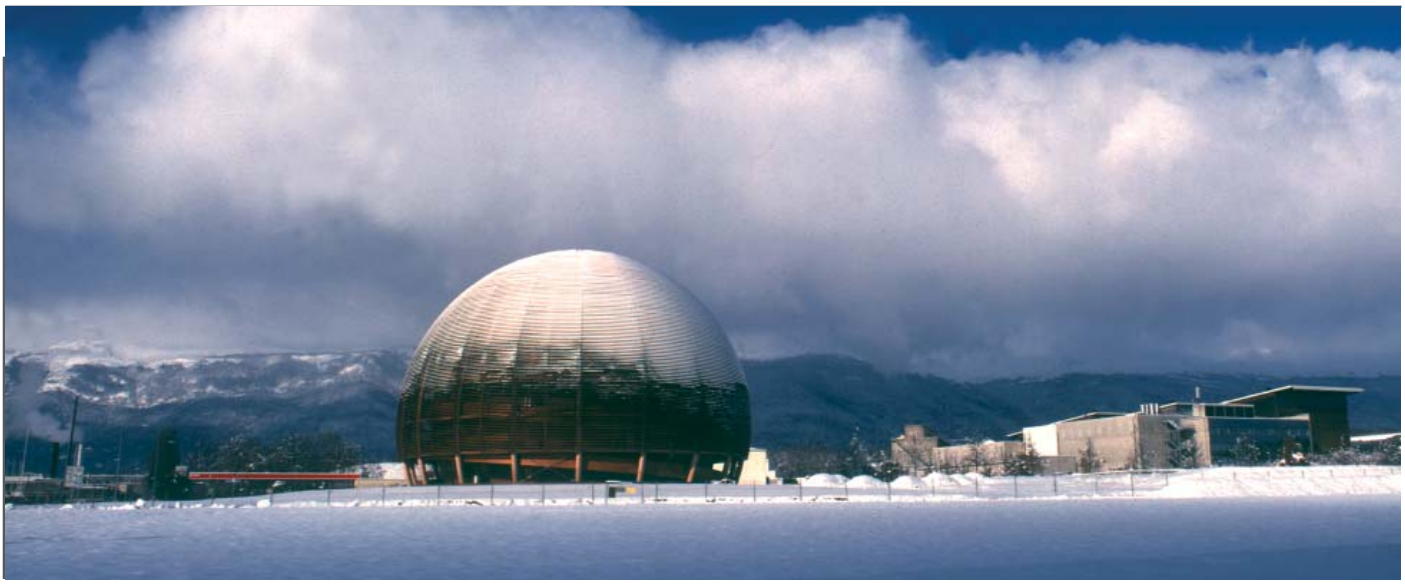

internally," explains Peter. "I think DataCentred has helped them sort out problems too. The technical experience of Matt Jarvis and Nick Jones, in particular, has helped us to share knowledge with these other UK institutions."

The collaborative approach has been a two-way street, says Peter. "It's great for us to work with Industry and to help Industry. We've put a large workload on the DataCentred systems and this can only be good for their customers. We've kept them busy and been a great test-bed for smoothing out problems in early operation. That's a good thing for us, for DataCentred, and for improving the infrastructure here in the North-West."

The computing demands of the ATLAS project will be greater in Run2 (2015-2017) than Run1 (2010-2012) and Peter expects cloud computing to play an increasing role during Run2.

The deployment at DataCentred is already on a par with Lancaster's computing grid resources. The next step will be to move some of the data storage to DataCentred.

DataCentred peers with Janet, the UK's academic network. "With high-bandwidth connectivity into that network," says Peter, "we can use DataCentred's storage efficiently. Being able to commission it quickly without buying hardware or employing sysadmins makes the DataCentred Ceph storage very attractive. In addition, we provide computing resources for a lot of different academic projects and, based on our experience with DataCentred, I can see cloud playing an increasing role there too – the advantages of cost and flexibility are very compelling."

"Our experience has been totally positive with DataCentred"

PETER LOVE

Peter says he has genuine praise for everyone at DataCentred: "Having a UK-based OpenStack cloud is ideal. Our experience has been totally positive with DataCentred. They are very technically knowledgeable. They are enthusiastic about working with us. They are full of good ideas about how we can work better. They are very responsive – being so responsive is a key part, the communication is quick, even out of hours. But the glaring thing is the stability of it. The stability of the DataCentred platform is a breath of fresh air – it just keeps on going."

PHONE: +44 (0)161 870 3981
EMAIL: enquiries@datacentred.co.uk
WEB: www.datacentred.co.uk

DataCentred