DELL EMC

intel

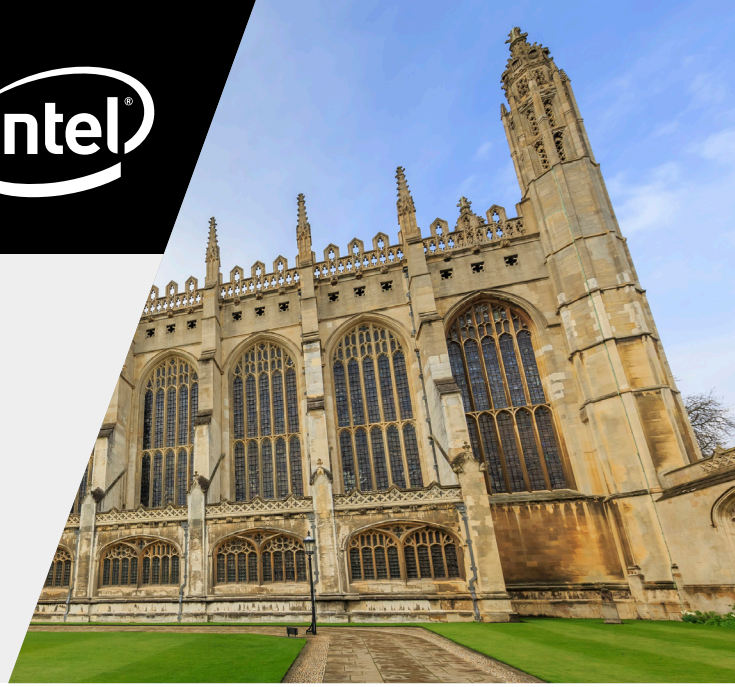

UK SCIENCE CLOUD

Researchers capitalise on the UK's largest academic supercomputer to expand the boundaries of artificial intelligence and data-driven science.

UNIVERSITY OF CAMBRIDGE

University of Cambridge | Research | United Kingdom

Business needs

The University of Cambridge needed a high-performance computing system to help solve some of today's most demanding data-driven simulation and AI challenges.

Solutions at a glance

- Dell EMC PowerEdge™ servers
- Intel® Xeon® processors
- Intel® Omni-path Architecture
- OpenStack software

Business results

- Enabling data-centric AI research for academic and industrial users
- Making HPC resources accessible to researchers via a cloud interface
- Advancing the UK government's industrial strategy

The Cumulus supercomputer provides performance rated at more than

2.27 petaflops

37,000 cores

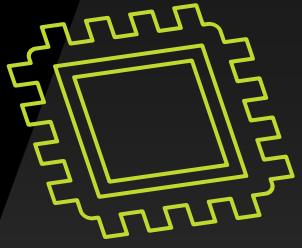

With the Dell Data Accelerator, the system can read data at

500 gigabytes/sec

2 million IOPS

The pursuit of excellence

The University of Cambridge is one of the world's oldest universities and most highly regarded centres of academic scholarship. Its reputation for outstanding academic achievement is known worldwide and reflects the intellectual achievement of its students and faculty, as well as the world-class research carried out at the University and its Colleges. That's all part of the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

These principles guide the work of the University's Research Computing Service, which provides leading-edge research computing services across all academic disciplines and to the broader realm of the UK scientific and industrial community. Today, to help its constituents stay at the leading edge of artificial intelligence (AI) and scientific research, the Research Computing Service has launched the UK's largest academic supercomputer — the Cumulus–UK Science Cloud.

The Cumulus–UK Science Cloud

The Cumulus system provides more than 2 petaflops of performance, powered by Dell EMC PowerEdge™ servers, Intel® Xeon® processors and the Intel® Omni-path Architecture (Intel® OPA). The system incorporates OpenStack software to control pools of compute, storage and networking resources and make them readily accessible to users via a cloud interface.

In addition to OpenStack-provisioned x86 and GPU bare-metal and virtualised hosts, Cumulus incorporates the highly innovative Dell Data Accelerator (DAC). This accelerator provides more than 500 GB/s of I/O read performance, which makes it the UK's fastest HPC I/O platform, according to the Research Computing Service. The result is a single heterogeneous x86/GPU platform that provides the UK's most advanced supercomputing cloud.

Viewed from a more granular level, the DAC system was built using Dell EMC PowerEdge™ R740xd 2U servers. Each DAC server has a 16-core Intel® Xeon® Gold CPU running at 2.5 GHz. Each of these servers contains a PCIe-x switch to connect six Intel® SSD P4600 Series NVMe solid state drives to each CPU, for a total of 12 SSDs for each server. The Cumulus installation uses 24 of these servers with 12 NVMe drives per server. To balance the system (based on assessing the peak performance of NVMe and network bandwidth), each of the 24 servers has two Intel OPA adapters.

In addition to the optimisation of the individual DAC server configurations, the Research Computing Service has worked closely with Dell EMC and Intel to optimise the network topology to best exploit the DAC configuration. And on the OpenStack side, the system makes use of the iDRAC capabilities embedded in Dell EMC servers for provisioning and management of the cluster, in particular with respect to Ansible-driven infrastructure as code.

The Cumulus system, which is open to academic and industrial users across the UK, was delivered in partnership with Dell EMC and StackHPC, a UK startup specialising in the convergence of HPC and cloud. It was funded with investments totalling more than £13 million from the UK's Engineering and Physical Sciences Research Council (EPSRC) Tier 2 HPC activity, the UK's Science and Technology Facilities Council (STFC) via the DiRAC and IRIS HPC programmes, and the University of Cambridge.

Empowering AI research teams

The Cumulus–UK Science Cloud has what it takes to help users solve extremely difficult data-driven, simulation and AI challenges, according to Dr. Paul Calleja, the University's Director of Research Computing Services.

"For people who need to do analytics or machine learning and process lots of data, we are bringing together on one system high levels of compute and high levels of I/O, combined with Hadoop and machine learning frameworks delivered within an OpenStack software environment, which allows both customisability and security for the tenants," Dr. Calleja says. "With all those things together, this machine can be used to deliver data-centric research to new and emerging communities."

Dr. Calleja notes that AI researchers require supercomputing capacity that is able to process huge amounts of data at very high speeds — which is exactly what Cumulus is designed to do.

The Cumulus–UK Science Cloud

is an outgrowth of the University of Cambridge's Peta4 supercomputer, launched in 2017. At 1.6 petaflops, Peta4 entered the Top500 supercomputer rankings as the fastest academic supercomputer in the UK and the 75th fastest system in the Top500 list.[1]

1 University of Cambridge Research Computing Services, https://www.hpc.cam.ac.uk/.

DELL EMC

"You cannot feed these people enough compute," he says. "They will eat whatever you give them. Cambridge's supercomputer provides researchers with the fast and affordable supercomputing power they need for AI work."

Many AI projects involving Cambridge researchers are already under way, according to Dr. Calleja.

"In the life sciences, we are working on medical imaging analysis and genomics, and in astronomy, scientists are using AI research to map exoplanets," he says.

The Square Kilometre Array project, or SKA is using OpenStack and the Data Accelerator to prototype next-generation HPC systems. The SKA is a new global radio telescope project, currently in the design study phase. The telescope is planned to be a "world array" that will be located partly in South Africa and partly in Australia. It will have sensitivity 100 times greater than the most sensitive radio telescopes of the present generation, and the ability to survey the sky up to 1 million times faster.[2]

"We also have some fairly groundbreaking work around genomics with the UK 10K project, where researchers are using the Hadoop data analytics engine to drive groundbreaking results around very large-population genome studies," Dr. Calleja says. Under that project, established in 2010, a consortium of biologists and genome scientists are working to facilitate the detailed gene variant analysis of 10,000 people to help elucidate the links between gene variation and illness.

Those are the kinds of challenges that are best conquered with the processing power of the world's Top500 supercomputers — like the Cumulus–UK Science Cloud.

"The cloud is open to all UK academic and industrial users," Dr. Calleja says, "and it is ready and able to solve the UK's most challenging data-driven, simulation and AI problems. Cumulus has what it takes to drive both pure research and industry product development — faster, further, bigger."

'Cracking the HPC storage problem'

Today, the University of Cambridge Research Computing Service is leveraging the Dell Data Accelerator (DAC) and the Distributed Name Space (DNE) feature in the Lustre file system to optimise the Cumulus cluster for top I/O performance. This optimisation work has led to a huge leap forward in storage performance, according to Dr. Paul Calleja, the University's Director of Research Computing Services.

In benchmark testing, the Cumulus system reached a score of 158 on the IO-500 list, he says. This score, which the Research Computing Services team has submitted to the list leaders, surpasses the current No. 1 score in the I/O 500 by more than 20 points, according to Dr. Calleja.

"With DNE, the IOPS performance of this solution is amazing," Dr. Calleja says. "The guys had to work around many Lustre bugs and adjust many Lustre parameters just to get it to run, but now we have stable, repeatable and very high performance runs with no error and determinant behaviour, so I think we have cracked the HPC storage problem."

2 University of Cambridge Department of Physics, Cavendish Laboratory, Square Kilometre Array (SKA) research.