

## HPC helps scientists tackle tomorrow's problems, today

University of Cambridge supercomputing resources enable breakthrough healthcare research with recent upgrade to 5X simulation capability and 20X AI performance.



### Business needs

The Cambridge Biomedical Campus is the largest centre of medical research and health science in Europe. On the campus, the School of Clinical Medicine is the third best medical school in the world, enabling doctors to pursue an academic career alongside clinical practice and there's an ever-growing need for High Performance Computing resources.

#### Solutions at a glance

- Dell PowerEdge servers
- Scientific OpenStack private cloud
- NVIDIA Mellanox QM8700 networking
- Lustre, NFS and object storage

#### Customer profile



**UNIVERSITY OF  
CAMBRIDGE**

Healthcare Research | United Kingdom



“We are answering the unmet need for secure, supercomputing resources using a breakthrough design that supports clinical medicine in a novel, fast and agile way to improve public and patient health.”

**Dr. Paul Calleja**

Director of Research Computing,  
University of Cambridge

### Business results

- Allows researchers to work with clinical medicine in a fast and agile way
- Revolutionizes how researchers can solve the world's biggest challenges
- 5X boost in simulation capability
- 20X increase in AI performance
- Reduces AI processing time from 2 years to one month
- Underpins collaboration for 50X larger, next generation exascale supercomputer

The prestigious University of Cambridge and its world-leading supercomputing resources is changing the way scientists tackle some of the world's most demanding research and health-related problems. Through a convergence of traditional simulation, artificial intelligence (AI), and data analytics, supercomputing enables next-level research with security, speed and agility. A key application area is clinical research to proactively improve healthcare outcomes and help solve tomorrow's problems, today.

Enabling world-class discovery and innovation, the University of Cambridge recently upgraded its powerful supercomputing resources. The upgrade generated a 5X boost in simulation capability and a 20X increase in AI performance to support thousands of projects. These projects allow researchers to work with large volumes of data sets processed through highly controlled, secure and complex computational pipelines to rapidly uncover new insights. In health-related areas such as COVID, cancer, and personalized medicine, supercomputing can greatly speed up discoveries, diagnoses and treatment plans.

Based in the UK, the University of Cambridge has some of the world's foremost supercomputers. The Cumulus system is built on more than 2,500 servers supporting more than 3,000 researchers in the UK and 700 world-leading projects. Peta 4 is comprised of compute nodes based on Dell PowerEdge servers with Intel® Xeon® Scalable processors. Wilkes3, one of the most powerful supercomputers for academic research in the UK, is comprised of PowerEdge XE8545 servers, each with 4x NVIDIA A100 GPUs. Wilkes3 achieved position 3 in the [June 2021 Green500 list](#), and position 100 in the [June 2021 Top500 list](#).

To stay on the cutting edge of technology, the University of Cambridge frequently upgrades these systems, collectively creating an innovative, heterogeneous cluster design using different types of computational resources in one platform, tailored to meet commercial and academic research computing needs. State-of-the-art Dell servers, networking and storage work together in a single system to optimize flexibility, usability, speed and security.

## Answering unmet need for secure compute power

During recent years, the volume of healthcare data has skyrocketed. Using and analyzing that data for insights enables next-level healthcare breakthroughs. However, that data also requires a unique security model to ensure proper handling of highly confidential patient data. Simultaneously, supercomputing resources must be employed to rapidly process large volumes of data and enable the sophisticated analytics required to extract meaningful insights.

The University of Cambridge's recent supercomputing upgrade is innovatively designed to manage these novel, multi-dimensional healthcare requirements. The technology runs on an ISO 27001 security domain within a flexible, open stack, allowing for specific security zones in one file. Previously, separate systems were required to achieve the necessary security levels.

"Our supercomputing resources keep data available while keeping it secure. Our upgrade design is a major breakthrough that allows us to support clinical medicine in a novel, fast, and agile way," states Dr. Paul Calleja, Director of Research Computing at the University of Cambridge. "With our technology upgrade, we are answering the unmet need for secure, supercomputing resources that effectively leverage rising volumes of confidential healthcare data with speed and ease to help improve patient and public health."

## Compute resources help manage COVID complexities

Spanning numerous areas, the University of Cambridge's supercomputing resources have become a vital tool for scientists and researchers, with a focus on healthcare and clinician collaboration. In one health-related example, those compute resources are enabling the Medical Research Council (MRC) Biostatistics Unit (BSU) to process and provide COVID insights on a weekly basis and better inform decision-making and COVID interventions.

Swift turnaround is essential. However, as COVID data sets grow, so does the complexity of the biostatistical models to analyze that data, in turn increasing the compute time needed to generate weekly updates. Supercomputing resources are required to handle this growing complexity and generate results in time so that policymakers can make better decisions regarding public health initiatives.



"The Cambridge Open Exascale Lab will make supercomputer resources widely available to exponentially catalyze powerful possibilities around the world."

**Dr. Paul Calleja**

Director of Research Computing,  
University of Cambridge

## Supercomputer pioneers personalized cancer treatments

Through an open source initiative called [OpenCB](#), the University of Cambridge is also delivering state-of-the-art clinical genomic interpretation, which allows clinicians to analyze and understand a patient's genome in relation to a specific disease.

In the case of ovarian cancers, for example, each cancer is often different in many ways, impacting each patient's response to treatment. By analyzing a patient's genome and then tailoring treatment based on sophisticated analysis and rich data sets of specific cancer types, clinicians are personalizing medicine and generating a higher probability of success. For cancer patients, speed is of the essence. In addition to enabling tailored treatments, the University of Cambridge's technology upgrades help to deliver a swift diagnosis and rapid treatment plan based on rich, data-driven insights.

The Computational Biology Lab of the [HPC Service](#) is focused on the development of new advanced computing solutions using the most modern HPC and big data technologies to solve current challenges in genomic data analysis and visualization. Some of these challenges include big data size, performance, scalability, security, data integration, collaboration, knowledge base, variant annotation and more.

## Co-designing with Dell for a flexible, heterogeneous environment

In the design and deployment of the University of Cambridge's supercomputer resources, collaboration between technology providers, service providers and users has created a unique advantage. "Our sixteen-year-long partnership with Dell has endured for several reasons. Dell has exactly the right model of technical and commercial fit to meet our needs.

Dell values and understands our co-design partnership. And, Dell has supported us and delivered value year after year. That's why we work with Dell," says Dr. Paul Calleja.

## What's next?

Never resting on its laurels, the University of Cambridge, in partnership with Dell and Intel, launched the [Cambridge Open Exascale Lab](#). At the Lab, the team focuses on next-generation supercomputer technology to deploy systems 50X larger than the ones used today. Exascale computing systems are capable of a billion billion calculations per second ( $10^{18}$ ), which will bring many opportunities and some challenges.

The Lab is focusing on areas such as power, programmability and middleware. "We are researching how you build, power and use exascale systems. A key rate-limiting factor is power consumption," shares Dr. Paul Calleja. "The Cambridge Open Exascale Lab is looking at technologies needed and how to democratize them. By making these resources available to everyone, we can exponentially catalyze powerful possibilities around the world."

### Additional resources:

#### Cambridge Open Exascale Lab

<https://www.exascale.hpc.cam.ac.uk>

#### Pioneering ovarian cancer research

<https://www.cam.ac.uk/stories/ovarian-cancer>

#### MRC BSU

<https://www.mrc-bsu.cam.ac.uk/about-bsu/>

#### Exascale Computing Project

<https://www.lanl.gov/projects/exascale-computing-project/>