

CLIMB PROJECT SUPPORTS RESEARCH COLLABORATION WITH RED HAT CEPH STORAGE



SOFTWARE AND SERVICES

Red Hat® Ceph Storage

Red Hat Consulting

HARDWARE

Dell EMC PowerEdge
R730XD Rack Servers

R730XD servers, with each server containing 16x 4TB HDDs, giving a total raw storage capacity of 6912TB. All data stored in this system is replicated 3x, which gives a usable capacity of 2304TB.

The Cloud Infrastructure for Microbial Bioinformatics (CLIMB) project is a collaboration between four UK universities to provide a comprehensive cloud-based platform for medical microbiologists in the UK to support rapid analysis and sharing of research data. To achieve this mission, CLIMB needed to create shared IT infrastructure to help researchers collaborate and produce new medical breakthroughs. The project chose to deploy RDO OpenStack® as a scalable, easily accessible cloud platform and Red Hat Ceph Storage to provide scalable object storage. With this solution, CLIMB can share a large amount of genome sequence data and project datasets between university sites and groups.



MEDICAL RESEARCH

BENEFITS

- Deployed scalable, flexible platform and object storage to support massive volumes of research data
- Enhanced collaboration through global access to tools and easier addition of new research institutions
- Reduced storage costs with flexible, cloud-ready solution

“Researchers can scale up resources as required. I’m excited by the number of publications coming out of the work we’re doing on OpenStack and Ceph, from simple biological analysis to more complex work.”

DR. THOMAS CONNOR
CO-INVESTIGATOR, CLIMB



facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

“We’ve had researchers from Gambia running analysis using the system, and we had teams sequencing data on OpenStack and Ceph last year during the Ebola outbreak in Africa. This is a large-scale, remote resource that we can ship data to, from anywhere.”

DR. THOMAS CONNOR
CO-INVESTIGATOR, CLIMB

SUPPORTING MEDICAL BREAKTHROUGHS THROUGH COLLABORATION

The Cloud Infrastructure for Microbial Bioinformatics (CLIMB) project is an IT research collaboration that involves four UK universities: Warwick, Birmingham, Cardiff, and Swansea Universities, with Bath University in the process of joining. Funded by the UK’s Medical Research Council (MRC), the seven-year project is closely monitored by the academic research and public health sectors. The project’s mission is to provide a comprehensive cloud-based platform for medical microbiologists in the UK to support rapid analysis and sharing of research data. Prior to CLIMB, researchers within UK universities would often buy, manage, and maintain local servers to meet their storage and analysis needs. As a result, institutions and research groups could not easily share software or reproduce data, key capabilities for scientific research. CLIMB’s goal is to use software-based solutions and commodity hardware to simplify collaboration.

“Genomic sequencing, for example, requires huge amounts of computing power,” said Dr. Thomas Connor, co-investigator for CLIMB. “However, most microbiologists are not computer experts. CLIMB provides a platform for IT resource access to help them quickly, easily work and collaborate.”

For example, diseases – such as HIV, Ebola, or hospital-acquired infections – are most effectively researched in collaboration with other experts. But isolated, fragmented IT systems made this critical collaboration difficult.

“Different universities running different systems made it very difficult to share data,” said Connor. “Collaboration was too challenging. Sharing resources is a great way of getting people to work together.”

DEPLOYING A STABLE YET FLEXIBLE CLOUD PLATFORM

CLIMB sought to create a robust yet easy-to-use platform that could scale on demand. In addition, the project needed storage for hundreds of terabytes of data that would scale to petabytes and beyond. To build this cloud platform, the project chose the RDO community version of OpenStack with Red Hat Ceph Storage for scalable object storage.

“It was an easy choice,” said Connor. “OpenStack provides the necessary stability and functionality. The more difficult decision was around storage, but it was clear Red Hat Ceph Storage was competitive on price and quality.”

Implementing the new solution was complex but made easier with help from Red Hat Consulting.

“It required a lot of effort to get going and to get it right,” said Connor. “But a large part of its success is due to Red Hat’s engagement. We had the same Red Hat consultant working with our IT team for a week at a time, providing great continuity. We now have a platform that fits our needs exactly.”

ACCELERATING INSIGHT WITH OPEN SOURCE SOLUTIONS

IMPROVED FLEXIBILITY AND SCALABILITY

The new solution, based on RDO OpenStack and Red Hat Ceph Storage, helps customers efficiently manage exponential data growth. CLIMB can now provide hundreds of individual virtual machines (VMs) supported by over 1.5PB of raw object storage per site. With RDO Openstack, CLIMB can provide VMs that can scale as needed, providing on-demand nodes ranging from single CPU cores and hundreds of MB of RAM to VMs with over 100 CPU cores and 3TB of RAM – considerably more than what research teams would normally need.

“Red Hat Ceph Storage provides object storage and block storage with incredibly high performance,” said Connor. “We can use it to store copies of images on our system, so that other researchers can download and use them. The features of Red Hat Ceph Storage let us share these large volumes efficiently. We’re confident that we have the storage volume we need to keep up with future generations of research data.”

EASIER COLLABORATION

With the Red Hat solution, the UK-focused project can expand globally. “We’ve had researchers from Gambia running analysis on the system, and we had teams uploading sequencing data to OpenStack and Ceph last year during the Ebola outbreak in Africa. This is a large-scale, remote resource that we can ship data to, from anywhere,” said Connor.

With the new solution, data can be shared more effectively among researchers. As a result, researchers react faster to infection outbreaks. In addition, improved collaboration helps researchers take full advantage of publicly funded grants through re-using and re-examining data – for example, to validate important or controversial claims.

“We’re good at applying lab work to real-world problems,” said Connor. “If you make infrastructure available, the research community will work together and take advantage of it. Microbiology has a particularly solid record of collaboration. Because we often focus on issues such as disease, we also tend to be looking at global problems.”

REDUCED STORAGE COSTS

Designed for the cloud, Red Hat Ceph Storage significantly lowers the cost of enterprise data storage by offering flexible, scalable storage that does not require additional hardware investment.

“The cost per terabyte of the CLIMB system is approximately half of the cost of pre-existing systems I’m involved with at Cardiff University,” said Connor.



FINDING NEW OPPORTUNITIES FOR INNOVATION

The new platform based on RDO OpenStack and Red Hat Ceph Storage has been quickly embraced by users.

“Researchers appreciate the freedom to enjoy the elasticity of the cloud,” said Connor. “I’m excited by the number of publications coming out of the work we’re doing, from simple biological analysis to more complex work,” said Connor.

The success of the Red Hat solution has prompted CLIMB to expand its use of Red Hat Ceph Storage. “We were so impressed that we’re now transferring some IBM GPFS [General Parallel Filesystem] storage we’d bought to Ceph as well,” said Connor.

CLIMB plans to use its enhanced capabilities to further support medical innovation by adding new universities to the project.

“We have an open system with an infrastructure that can support other institutions. We’re actively looking for others to join,” said Connor. “We’re also looking at extending the framework to other sectors.”


ABOUT CLIMB

The CLIMB project (Cloud Infrastructure for Microbial Bioinformatics) is a collaboration between Warwick, Birmingham, Cardiff, and Swansea Universities to develop and deploy a world-leading cyber-infrastructure for microbial bioinformatics; providing free cloud-based compute, storage, and analysis tools for academic microbiologists in the UK.

The project is supported by three world-class Medical Research Fellows, a comprehensive training program, and two newly refurbished bioinformatics facilities in both Warwick and Swansea. CLIMB is funded by the UK’s Medical Research Council.

ABOUT RED HAT

Red Hat is the world’s leading provider of open source software solutions, using a community-powered approach to provide reliable and high-performing cloud, Linux, middleware, storage, and virtualization technologies. Red Hat also offers award-winning support, training, and consulting services. As a connective hub in a global network of enterprises, partners, and open source communities, Red Hat helps create relevant, innovative technologies that liberate resources for growth and prepare customers for the future of IT.



facebook.com/redhatinc
@redhatnews

linkedin.com/company/red-hat

NORTH AMERICA
1 888 REDHAT1

**EUROPE, MIDDLE EAST,
AND AFRICA**
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com

The OpenStack® Word Mark and OpenStack Logo are either registered trademarks / service marks or trademarks / service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation’s permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community.