


# CORNELL UNIVERSITY CHOOSES RED HAT STORAGE FOR DEMANDING HIGH-PERFORMANCE COMPUTING ENVIRONMENT

Cornell University's Center for Advanced Computing (CAC) required a highly scalable and reliable storage solution for its clients' data-intensive research projects in high performance computing environments. The Red Hat Storage Software Appliance allowed the CAC to meet the needs of a key constituent, Cornell's Institute for Biotechnologies and Life Science Technologies, which needed a storage platform robust enough to perform DNA sequencing, proteomics and imaging.



Cornell University

## RED HAT SOFTWARE

- ▶ Red Hat Storage Software Appliance

CUSTOMER SINCE  
**2010**

HIGHER EDUCATION  
INDUSTRY

*"With Red Hat Storage, we were able to dramatically avoid expenditures with a low-cost software solution, while keeping our current infrastructure in place. It enabled us to scale easily and affordably."*

JAMES VANEY, IT DIRECTOR FOR THE INSTITUTE  
FOR BIOTECHNOLOGY AND LIFE SCIENCE  
TECHNOLOGIES AT CORNELL UNIVERSITY



LOCATION:  
Ithaca, New York

---

**HARDWARE**

158 terabytes using Dell R610  
connected via InfiniBand to the  
DataDirect Networks  
(DDN) 9900

**OPERATING SYSTEM**

Mix of native OS and Windows

---

**CORNELL UNIVERSITY: CUTTING-EDGE RESEARCH ON HIGH-PERFORMANCE COMPUTING**

Founded in 1865, Cornell University is the largest university in the Ivy League, with more than 20,000 students from every state and 120 countries around the world. Offering a unique combination of renowned scholarship and democratic ideals, Cornell mixes practical subjects with the classics and is the most educationally diverse member of the Ivy League.

The Cornell University Center for Advanced Computing (CAC) was founded in 2007 to help the university with high-performance computing initiatives to support data-intensive research efforts across the university, government, and industry. An early adopter of emerging technology, CAC enables scientists around the world from a variety of disciplines to accelerate research success. CAC operates 15 Red Hat® Enterprise Linux® high-performance computing clusters, and is known for its expertise in analyzing the usability of large-scale computational systems and for developing, managing, and evaluating training and education programs.

**SEEKING A STORAGE SOLUTION TO PROMOTE RESEARCH, EDUCATION, AND TECHNOLOGY TRANSFER IN BIOTECHNOLOGY**

CAC has many constituents across academia, government, and industry. One of these is Cornell's Institute for Biotechnology and Life Science Technologies. James VanEe, IT director for the Institute, came to the technical staff at CAC in 2010 for help identifying and deploying a storage solution that would help his department with biotechnology research that ultimately benefits the environment, agriculture, engineering, and veterinary and human medicine.

Because the Institute for Biotechnology and Life Science Technologies brings together university scientists conducting research in biology and the physical, engineering, and computational sciences, "the collaborative nature of the Institute and its research produces extremely large amounts of data," said Steven Lee, CAC systems consultant.

At the time, the Institute's scientists couldn't access unstructured data and were limited by standard file systems that capped at 8 terabytes and 16 terabytes per node. Since the Institute produces more than 20 terabytes of data a month, it needed a solution that could both scale and was highly available. The scientists also required a file system that didn't have the access limits of typical file systems, so a global namespace was a necessity.

"The idea of a scale-out storage solution was something we'd always been interested in, but never could implement due to cost," said VanEe. He had considered other solutions, including Isilon, but the large capital cost of such systems was prohibitive. "So we turned to CAC for suggestions on other possible storage solutions that did not require a big upfront investment."

## RED HAT STORAGE BEST SOLUTION FOR CAC'S HIGH-PERFORMANCE COMPUTING CLUSTERS

VanEe first heard about Red Hat Storage, formerly Gluster, at Bio-IT World. He was immediately impressed with the software-only feature of the product, as that would allow him to leverage CAC's existing storage hardware infrastructure. He also found attractive the scale-out architecture and the large global namespace that would ensure high performance and availability for even the largest data sets.

After CAC analyzed several vendor offerings, the Institute chose Red Hat Storage Software Appliance as the best solution. It deployed the solution on a mix of native environments, on 158 terabytes of storage using a Dell PowerEdge R610 rack server connected via InfiniBand to a DDN 9900 data infrastructure platform.

The installation of the Red Hat Storage Software Appliance was easily completed within a few days, and has been running flawlessly in production for more than a year. Today, the Institute is primarily using Red Hat Storage to store the data generated by its DNA sequencing research. It is also using it for archival copies of data from other groups within the department.



### GREATER SCALABILITY AND AVAILABILITY AT A LOWER COST

With Red Hat Storage Software Appliance, Cornell's Institute for Biotechnology and Life Science Technologies was able to continue utilizing its existing storage disks. This enabled the Institute to scale affordably using a high-performance SAN. In addition to not having to replace its installed base, the Red Hat software-only solution allowed the Institute to scale as much as it needed without deploying additional servers and storage hardware.

"Cost avoidance was one of our primary goals when looking for a new storage solution," said VanEe. "We simply did not have the budget available for an upfront capital investment in equipment. With Red Hat Storage, we were able to dramatically avoid expenditures with a low-cost software solution, while keeping our current infrastructure in place. It enabled us to scale easily and affordably."

Additionally, Cornell can now provide scientists access to data without being limited to standard file systems. Red Hat Storage allows them to work more easily with unstructured data. "Selecting a solution with a global namespace allowed us to eliminate administrative and data management overhead, saving a significant amount of time and cost," noted VanEe.

With the Red Hat Storage Software Appliance, the Institute can continue growing its research programs, while increasing researcher productivity. "One of my main goals as IT director is to create an environment where new technologies can be adopted at the drop of a hat," said VanEe. "If our department spends half a million dollars on a new sequencer, they don't want to wait to store the data. Red Hat Storage helps us stay ahead of the curve with its flexibility to fit in with new technologies."



---

## ABOUT RED HAT

Red Hat was founded in 1993 and is headquartered in Raleigh, NC. Today, with more than 70 offices around the world, Red Hat is the largest publicly traded technology company fully committed to open source. That commitment has paid off over time, for us and our customers, proving the value of open source software and establishing a viable business model built around the open source way.

## SALES AND INQUIRIES

**NORTH AMERICA**  
1-888-REDHAT1  
[www.redhat.com](http://www.redhat.com)

**EUROPE, MIDDLE EAST  
AND AFRICA**  
00800 7334 2835  
[www.europe.redhat.com](http://www.europe.redhat.com)  
[europe@redhat.com](mailto:europe@redhat.com)

**ASIA PACIFIC**  
+65 6490 4200  
[www.apac.redhat.com](http://www.apac.redhat.com)  
[apac@redhat.com](mailto:apac@redhat.com)

**LATIN AMERICA**  
+54 11 4329 7300  
[www.latam.redhat.com](http://www.latam.redhat.com)  
[info-latam@redhat.com](mailto:info-latam@redhat.com)