




# IVY LEAGUE UNIVERSITY MANAGES DATA-INTENSIVE RESEARCH WITH RED HAT STORAGE

Cornell University searched for a highly reliable, available, and scalable storage solution to help manage data-intensive research projects, such as DNA sequencing. The flexibility and reliability of Red Hat Storage has enabled the Institute to achieve the growth needed to continue its research programs.



Cornell University

## RED HAT SOFTWARE

- ▶ Red Hat Storage

CUSTOMER SINCE  
**2010**

**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 VANEE  
CORNELL UNIVERSITY

LOCATION:  
North America



**HARDWARE:**

Mix of native OS and Windows on 158 terabytes using Dell R610 connected via InfiniBand to the DataDirect Networks (DDN) 9900

**A STORAGE SOLUTION TO PROMOTE RESEARCH, EDUCATION, AND TECHNOLOGY**

The Cornell University Center for Advanced Computing (CAC) conducts research on current and emerging technologies to enhance the capabilities of high-performance computing. From system firsts to the development of new technologies, the center rapidly adopts and extends advanced information technologies in order to solve its clients' most pressing problems.

James VanEe, IT Director of Cornell's Institute for Biotechnology and Life Science Technologies, looked to the technical staff at CAC to identify and deploy a storage solution for its department. VanEe was looking for a solution that promoted research, education, and technology transfer for applications of biotechnology that ultimately benefit the environment, agriculture, engineering, and veterinary and human medicine.



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

"The storage solution the Institute was using at the time only provided access to standard file systems and capped at 8 and 16 terabytes per node; they clearly needed a storage solution that would allow them to have seamless access to all data in every node. They also required a file system that didn't have

the data access limits of typical file systems. In the end, they knew that a global namespace was a necessity," explained Lee.

**RED HAT STORAGE PROVIDES ELASTIC SCALING CAPABILITIES**

VanEe was first introduced to Red Hat Storage, formally Gluster, at Bio-IT World. He was initially attracted to the software-only feature of the product, since it could quickly add value to Cornell's existing infrastructure. Other features that stood out were the solution's scale-out architecture and the large global namespace that would ensure high performance and availability.

"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. "We considered other solutions—such as Isilon—but the large, up-front capital investment motivated us to turn to CAC for suggestions on other possible storage solutions that did not require a big upfront investment."

Since the Institute produces over 15 to 20 terabytes of data a month, this required a solution that could provide elastic scaling capabilities, remain highly available, and still be capable of handling large amounts of data output at any given time. After CAC analyzed several vendor offerings, the Institute opted to deploy Red Hat Storage on a mix of native and Windows environments, on 158 terabytes using Dell R610 connected via InfiniBand to a DDN 9900 data infrastructure platform.

The technology was up and running within a few days. Today, the department 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. The product has been running flawlessly in production for over a year now.

### RED HAT STORAGE HAS THE FLEXIBILITY TO KEEP CORNELL AHEAD OF THE CURVE

With Red Hat Storage, Cornell's Institute for Biotechnology and Life Science Technologies was able to lay new technology on its already existing disks, allowing for affordable scalability using a high-performance SAN. The organization not only saved money with Red Hat Storage, but avoided the potential high cost of 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 required a storage solution with a global namespace to provide access to data without the limits of typical file systems. Red Hat Storage removed the constraint of trying to keep the data structured. "Selecting a solution with a global namespace allowed us to eliminate administrative and data management overhead, saving us a significant amount of time and cost," noted VanEe.



The flexibility and reliability of Red Hat Storage allowed the Institute to achieve the growth needed to continue its research programs, while increasing researcher productivity due to the high availability of the data. With the elastic scaling capabilities provided by Red Hat Storage, the Institute relieved the constraint and pain of trying to manage unstructured data.

“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. We also appreciate Red Hat’s great support and look forward to continuing our excellent relationship.”

### About Cornell University

*Founded in 1865, Cornell University is the largest university in the Ivy League with over 20,000 students from around the world. Located in the heart of New York’s Finger Lakes region, Cornell offers more than 4,000 courses, 70 undergraduate majors, 93 graduate fields of study, undergraduate and advanced degrees, as well as continuing education and outreach programs. Cornell also supports more than 100 interdisciplinary research organizations, bringing faculty and students together to pursue research. The university has three national research centers serving broad scientific communities and reflecting a partnership of academia, government, and industry.*

---

#### 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

<b>NORTH AMERICA</b> 1-888-REDHAT1 www.redhat.com	<b>EUROPE, MIDDLE EAST AND AFRICA</b> 00800 7334 2835 www.europe.redhat.com europe@redhat.com	<b>ASIA PACIFIC</b> +65 6490 4200 www.apac.redhat.com apac@redhat.com	<b>LATIN AMERICA</b> +54 11 4329 7300 latammktg@redhat.com
---	--	--	--