

GAASTERLAND LABORATORY

THE CUSTOMER

The Gaasterland Laboratory of Computational Genomics at The Rockefeller University aims to create and use new software tools to explore the unfolding world of genomes. The tools are designed to integrate, analyze, and visualize the output of high-throughput molecular biology experiments in the context of complete genome sequence data. In an effort to apply new tools to specific biological questions in the most efficient way possible, a prerequisite is a high performance storage solution that is easy to access and manage.

THE CHALLENGE

When the Gaasterland Laboratory started in 1998, Dr. Terry Gaasterland, Associate Professor of The Rockefeller University, initially speculated that she would only need 500 GB of data capacity. However, as the files to be analyzed grew in number and size, she quickly realized that she needed far more storage. “Our data needs were huge,” said Dr. Gaasterland. “The data sets were growing and changing all the time.” Not only were the number of data sets to test increasing, but as new tools were developed, each tool would result in a new data set that needed to be stored and accessed for comparison to the original. Data growth was explosive and managing these ever-changing data sets was an administration burden.

To attack this problem, Dr. Gaasterland set out to find a storage system to achieve three key objectives. First, the system had to scale seamlessly without disruption to current users or mount points. Second, it needed to provide fast data throughput in order to reduce the time necessary to analyze data. Genomic data analysis is I/O bound and the new storage solution should eliminate the I/O bottleneck traditionally experienced with network storage solutions. Finally, the system needed to be extremely reliable and available to an extended community outside the Lab.



Industry: Life Sciences

The Challenge:

Explosive data storage had caused I/O bottlenecks, increasing the time required to develop new software tools used in molecular biology experiments. Managing the storage system had become a huge administrative burden and it was disrupting productivity, increasing analysis time and posing data availability problems.

Panasas Solution:

The fully integrated software/hardware solution included the Panasas® ActiveScale® Operating Environment and the PanFS™ parallel file system with the Panasas DirectFLOW® protocol.

Key Results:

- Improved System Reliability
- Increased I/O Performance
- Exceptional Ease of management

THE SOLUTION

During the evaluation process, Dr. Gaasterland and her team entertained solutions from IBM, Sun and Apple, but the Panasas® Parallel Storage Cluster was selected for its performance, reliability and ease of management. Having experience with distributed file systems, Dr. Gaasterland knew that the Panasas ActiveScale

“The system is blazingly fast. We’ve been able to eliminate our I/O bottleneck and thus carry out data analysis more quickly.”

- Terry Gaasterland
Associate Professor of
The Rockefeller University



Operating Environment architecturally had the capability to solve all their problems. The multi-TB solution was deployed to interact with Apple and Sun desktop workstations.

“Panasas met all the requirements on my checklist,” said Dr. Gaasterland. “Even though Panasas is a new company, we felt confident that they could provide not only the product but the support necessary to meet all our needs.”

THE VALUE

In the six months after moving into production, the Panasas Parallel Storage Cluster was everything the team expected and more. “The system is blazingly fast,” said Gaasterland. “We’ve been able to eliminate our I/O bottleneck and thus carry out data analysis more quickly.” The Panasas system is designed from the ground up to deliver exceptional random I/O and data access throughput.

Further, the single seamless namespace and dynamic growth capabilities had a significant impact on the organization. “The product is ‘plug-and-play’ at all levels,” commented Gaasterland. “Our team doesn’t have to think about it at all. We can focus our attention on biology research questions instead of systems administration

issues.” With the Panasas Storage Cluster, there is no reconfiguration required to expand data sets. The manageability features enable the system to grow without having to grow the overhead to support it.

Finally, during a recent two-day power outage the system’s reliability was tested and proved to be more than capable. While other vendors’ systems failed to provide a graceful shutdown to guarantee data integrity, the Panasas Storage Cluster was the only storage system in the lab that maintained complete data integrity. “At the end of the day, data availability and integrity is essential to do our job,” Gaasterland said. “Panasas proved itself to be more than reliable enough for our environment.”

“Our team doesn’t have to think about it at all. We can focus our attention on biology research questions instead of systems administration issues.”

- Terry Gaasterland
Associate Professor of
The Rockefeller University



Accelerating Time to Results™

6520 Kaiser Drive Fremont, California 94555 Phone: 1-888-PANASAS Fax: 510-608-4798 www.panasas.com
1-888-PANASAS (US & Canada) 00 (800) PANASAS2 (UK & France) 00 (800) 787-702 (Italy) +001 (510) 608-7790 (All Other Countries)

©2007 Panasas Incorporated. All rights reserved. Panasas, the Panasas logo, Accelerating Time to Results and ActiveScale are trademarks or registered trademarks of Panasas, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.