

## J. Craig Venter Institute

Protecting Data and Streamlining Storage for World-Class Research with Symantec Solutions



**The J. Craig Venter Institute (JCVI), a world leader in genomic research, had a variety of hardware and operating systems and needed a backup platform that could accommodate them. By standardizing on Veritas NetBackup, JCVI has been able to save thousands of dollars by repurposing a backup server, achieve a 99 percent data recovery success rate, and project a 100 percent backup success rate once tape backups are eliminated. JCVI can now allocate and change storage on the fly, without downtime, using Veritas Storage Foundation.**

### ORGANIZATION PROFILE

The J. Craig Venter Institute ([www.jcvi.org](http://www.jcvi.org)) is a non-profit genomic research institute with more than 400 scientists and staff, over 250,000 square feet of laboratory space, and locations in Rockville, Maryland and La Jolla, California.

### INDUSTRY

Life Sciences (Genomic Research)

### SOLUTION

Backup and Recovery; Disaster Recovery; Data Center Automation; Storage Management

### Diverse Genes, Diverse Technology

Consider the complexity of just one of the projects that the J. Craig Venter Institute is working on--the human genome. In 2007 the JCVI team sequenced the first complete diploid human genome, generating 20 billion base pairs of DNA. This project required several terabytes of storage space, and is just one of the many genomics research projects that the J. Craig Venter Institute (JCVI) is engaged. Genomic research is the key to understanding inter- and intra-species diversity and evolution. The JCVI, a world leader in genomic research, reflects this diversity in its workforce—and its technology environment.

JCVI uses many major operating systems, including Linux, Solaris, Windows and Mac OS X, running on a heterogeneous mix of hardware. And although the institute has only 400 employees, due to its complex research, it has 200 terabytes of data—a volume that exceeds that of many Fortune 500 corporations. Clearly, JCVI requires a backup solution that can accommodate its large data stores and work well with a variety of hardware and software.

### Wanted: Hardware Agnostic Backup Solution

In October 2006 several legacy organizations including The Institute for Genomic Research (TIGR), were merged to form The J. Craig Venter Institute. At the time, the newly formed organization's backup infrastructure was split between Veritas NetBackup™ and EMC Legato NetWorker.

“At TIGR, where I worked, we had been a Veritas NetBackup user since 2002, and I was already comfortable with the product,” says Eddy Navarro, Storage Team Lead, JCVI. “But I wanted to keep an open mind, so we did some research and it became clear to us that NetBackup was the logical way to go. It's a forward-looking product, especially considering its OpenStorage API. A feature like this that allows the software to be vendor agnostic is very important to us, as we did not want to be tied into a specific hardware vendor.”

**“We've never had any problems with Veritas NetBackup. The way the software is structured with policies and schedules makes it very clear and easy to use.”**

**Eddy Navarro**

Storage Team Lead  
J. Craig Venter Institute

Using Veritas NetBackup, The J. Craig Venter Institute has maintained a data recovery success rate of 99%.

**“The Veritas products have always come through for us in situations where our users needed a critical restore or uninterrupted access to genomic data.”**

**Eddy Navarro**

Storage Team Lead  
J. Craig Venter Institute

The organization is also in the process of converting from a disk-to-tape approach for its weekly full backups to a more efficient disk-to-disk environment, so it needed a backup solution that would work with a variety of virtual tape libraries (VTLs).

“It seems like the majority of the VTL vendors out there are working with Symantec first and foremost to integrate their products with NetBackup,” says Navarro. “NetBackup just made sense for us.”

By standardizing on NetBackup and centralizing backup administration, JCVI freed up an enterprise class server as well as IT resources that were previously dedicated to NetWorker.

### 99 Percent Successful Data Recovery

Navarro estimates that JCVI’s data recovery success rates have been in the 99 percent range since standardizing on NetBackup. “Our backup success rates are somewhat lower than that, but that’s not NetBackup’s fault,” he says. “It’s just that we have so much data, we’re running into space limitations with our tape drives. Once we start doing disk-to-disk backups, I expect our success rate will be 100 percent.”

With daily change rates for incremental backups approaching seven terabytes, JCVI needs comprehensive data protection for its database. The Veritas NetBackup for Microsoft SQL Server Agent provides an intuitive interface to manage consistent policies, and allows the institute to keep its databases online during backup and recovery, ensuring high availability.

Navarro is also happy with NetBackup’s ability to perform fast, image-based backup and recoveries of VMware virtual machines, and is looking forward to new functionality in the upcoming release of NetBackup that will increase the concurrency of multiple virtual machine backups.

JCVI called on Symantec Essential Support Services when it came time to upgrade to NetBackup 6.5. “The upgrade went seamlessly,” says Navarro. “They helped us make sure we didn’t have any catalog inconsistencies or other problems.” Symantec Educa-

## SOLUTION AT A GLANCE

### Business Drivers

- Enhance reliability of data backup and recovery
- Maximize availability of genomic data for scientists

### Technology Challenges

- Support a heterogeneous mix of hardware and operating systems
- Accommodate high data volume and change rate
- Simplify storage management and migration while maximizing uptime

### Solution

Symantec backup, recovery, and storage management for a heterogeneous infrastructure

### Symantec Products

- Veritas NetBackup™ 6.5, including the Microsoft SQL Server Agent
- Veritas Storage Foundation™
- Veritas Storage Foundation for Windows

### Symantec Services

- Symantec Education Services
- Symantec Essential Support Services

### Technology Environment

- Applications: Custom applications for genomic analysis, Microsoft Exchange Server 2003, VMware ESX Server
- Databases: Sybase, Oracle, Microsoft SQL Server
- Servers: A variety of servers and operating systems, including Red Hat and SUSE Linux, Solaris, Microsoft Windows Server 2003 and Mac OS X Server
- Storage: NetApp NAS, EMC Symmetrix and CLARiiON

## BUSINESS VALUE AND TECHNICAL BENEFITS

- Backup solution compatible with a variety of hardware, including virtual tape libraries
- Saved thousands of dollars and IT administration time by repurposing a server dedicated to former backup solution
- 99% successful data recovery
- Projected 100% backup success rates once tape backups are eliminated
- Fast, image-based backup and restore of virtual machines
- Ability to partition and grow storage with no downtime
- 48-hour storage downtime avoided through ability to migrate data on the fly

tion Services also provided training on the new version to two of JCVI's storage administrators.

"We've never had any problems with Veritas NetBackup," says Navarro. "The way the software is structured with policies and schedules makes it very clear and easy to use."

### Managing Data Growth With Zero Downtime

JCVI is also using Veritas Storage Foundation for online heterogeneous management of both its Solaris and Microsoft Windows-based storage arrays. Based on Veritas Volume Manager and Veritas File System, Storage Foundation provides a set of integrated tools to centrally manage data growth on the fly, without taking systems offline.

"We wanted easy expandability, maximum availability, and smooth data migrations, and Storage Foundation has performed quite well for us," says Navarro. "For example, when we need to migrate data between storage arrays, Storage Foundation allows us to have a virtualized storage layer on the host. This way, we can present the storage to the host, do a copy within

Volume Manager, and seamlessly migrate the data—as well as grow the size of the underlying storage—with no disruption to data availability."

Navarro describes a recent situation in which he used Storage Foundation to migrate Sybase data from one storage array to another.

"It completed in a matter of hours with no downtime, and that is the key," he says. "If I had to do this manually, we would be looking at a two- to three-day outage, which obviously would have been a major disruption to our users. It's hard to attach a dollar figure to that downtime, but 400 scientists and staff without adequate access to data for three days would have a major impact on our research."

### Enabling an Effective and Reliable Data Center

If datacenters had genetic traits, easy backups and restores and streamlined storage management would certainly be among the most prized. Navarro is confident that JCVI's Symantec software provides these advantages.

"The Veritas products have always come through for us in situations where our users needed a critical recovery or uninterrupted access to genomic data," he says.

<sup>1</sup> Human Genome Project Information: While JCVI is not endorsing or recommending this product, because of our positive experience with Symantec NetBackup and Storage Foundation, we are happy to let others know our opinion on how it has helped us solve some of our IT challenges.