

Compaq Tru64 UNIX
on AlphaServer Systems



Genentech discovers *high availability can be simple*

*Biotechnology pioneer
depends on Compaq
Tru64 UNIX with
TruCluster Server on
AlphaServer platform
for continued growth*

*Four-node cluster
supports more
than a terabyte of
StorageWorks storage*

“This stuff is seriously cool. Adding a new node into the cluster while online takes less than 30 minutes — and the additional management overhead is practically nil.”

Scooter Morris, Principal Systems Architect, Computing Technologies Group, Genentech

Back in 1976, when venture capitalist Robert A. Swanson and biochemist Dr. Herbert W. Boyer founded Genentech, Inc. (NYSE: DNA), the idea of using human genetic information to develop and manufacture new kinds of medicines was met with skepticism from both the academic and business communities.


But Genentech quickly proved their detractors wrong. Today the company is recognized as a pioneer in the field of biotechnology. And it has been making science, industry and business news ever since. In 1977, Genentech produced the first human protein (somatostatin) in a microorganism (E. coli bacteria). In 1982, it produced the first pharmaceutical based on DNA technology: recombinant human insulin, licensed to another company. And in 1985,

Genentech was the first company to shepherd a bioengineered medicine — one used for treating children with growth hormone deficiency — all the way from research, through manufacturing, to market.

Today, Genentech is responsible for more than a dozen of the approved biotechnology products currently on the market.

“Our research is primarily directed towards the development of therapeutics,” says Scooter Morris, principal systems architect, and head of the Computing Technologies Group at Genentech. The company focuses on developing medicines in the areas of cardiovascular medicine and cancer biology, as well as for specific diseases such as growth hormone deficiency, cystic fibrosis and asthma.

COMPAQ
NonStop™



Morris and his organization are responsible for ensuring that the computing infrastructure is in place to support Genentech scientists, its other employees, contractors and partners in developing, manufacturing and bringing these new products to market.

“Genentech’s drug discovery process relies heavily on computing resources, as does most other aspects of our business,” says Morris. “Our world-class computing infrastructure enables us to respond to changing automation needs as well as to enhance the productivity of every employee.”

Tru64 UNIX on AlphaServer provides the foundation

Genentech has long relied on *Compaq AlphaServer* systems running the *Tru64 UNIX* operating system with *StorageWorks* systems as a key element in its IT infrastructure. Applications running on the company’s *AlphaServer* systems include Netscape Enterprise Server, Oracle, Informix, Lawson, sendmail, Math software, High Performance Computing applications for research and many in-house, custom-developed applications.

Now in anticipation of continued dramatic growth over the next five years, Genentech is basing its entire e-mail and Web infrastructure on a clustered *AlphaServer* platform running *Tru64 UNIX* and *TruCluster Server* software with *StorageWorks* systems.

“The new cluster is built around four clustered *AlphaServer ES40* systems, *StorageWorks Fibre Channel RAID Arrays* and *Tru64 UNIX* and *TruCluster Server* software V5,” says Morris.

“We are very enthusiastic about the increased scalability and significantly decreased management costs that are a part of the V5 release. Not only do we see a substantial performance improvement, we do so with only one third of the floor space and with a dramatic reduction in complexity and an associated increase in manageability.”



▲ AlphaServer ES40



Meeting user demands for superior technology

Approximately 4,500 Genentech employees and contractors access email, file and Web services through their PCs, Macs and workstations.

With a dozen new products in clinical testing — and for Genentech scientists actively engaged research activities such as Bioinformatics, Protein Engineering and Molecular Biology — downtime is not an option.

“Genentech scientists are the most prolific in the biotechnology industry,” says Morris. Indeed, they publish at a rate of about 250 scientific papers a year — or more than one paper for every business day!

The company also leads in the highest average number of citations by other authors per published scientific paper, a peer-related measure used to indicate scientific impact. And Genentech holds more than 3,200 patents worldwide — with 2,000 more pending.

High-performance, highly available IT services are critical to helping Genentech attract and keep the industry’s top people. And people, says Morris, “are the major competitive advantage Genentech has.”

Bullet-proof availability to support even the most business-critical applications

Tru64 UNIX with *TruCluster Server* software on the *AlphaServer* platform delivers the performance and bullet-proof availability that Genentech needs to power the business processes that will keep it in the lead in a fast-changing and dynamic industry.

With delay of a product approval costing as much as \$1 million per day, any loss of productivity in Genentech’s product development or research efforts is unacceptable. The company’s e-mail system, which is used both for internal communications and collaborative communications with business partners and scientific collaborators, is considered business-critical with a goal of 100% uptime.

“Downtime is extremely costly both in terms of loss of productivity and in terms of opportunity costs,” says Morris. “*Tru64 UNIX* with *TruCluster Server* software allows Genentech to plan downtime for individual system nodes without downtime for the users. The common cluster address and fast failover of both storage and compute capacity make this task much easier and much more transparent.”

The rolling upgrade capability of the *TruCluster Server* system enables Genentech to add new servers and storage, to install maintenance patches and upgrade to new versions of operating and applications software — while applications keep running on at least one system in the cluster the whole time.



Performance for today's — and tomorrow's — applications

Rapid and dramatic growth is a fact of life at Genentech. In 1992, the company opened the world's largest biotechnology research facility, an \$85 million, 275,000-square-foot facility overlooking San Francisco Bay. A \$62.5 million Process Science Center was completed in 1995. More recently, the company opened the world's largest biotech manufacturing facility, a \$250 million, 310,000-square-foot facility in Vacaville, California. In all, Genentech occupies more than 1.9 million square feet of space in 21 buildings in South San Francisco and 325,000 square feet of space in Vacaville.

Supporting this kind of growth, says Morris, makes the easy scalability of both processing power and storage critical.

Tru64 UNIX is a high-performance, advanced enterprise-strength operating system offering scalability to multi-terabyte configurations including up to a 16 TB file system and 28 GB of memory, millions of concurrent users, more CPU power, dynamic tuning,

Web-based system management, and partitioning. No other UNIX solution offers as much scalability as *Tru64 UNIX* on *AlphaServer* solutions.

The Genentech configuration clusters four 4-way SMP *AlphaServer ES40* servers with 4 GB memory each using the 100 MB/sec MC2 Memory Channel cluster interconnect. Attached to the cluster is more than 2 TB of *StorageWorks Fibre Channel* storage.

"We chose the *StorageWorks* platform for its tight integration with *Tru64 UNIX* and *TruCluster Server* software," says Morris.

Growing with ease

Because *Tru64 UNIX* with *TruCluster Server* software on *AlphaServer* solutions scale up easily, Genentech can invest in the solution that meets their needs today, confident that the power and capacity will be there to support tomorrow's growth.

With capacity for four billion User Identifications (UID)/Group Identifications (GID), *Tru64 UNIX* solutions provide plenty of headroom for future growth.

Compared to the 1 to 2 TB of storage supported on other UNIX platforms, *Tru64 UNIX* supports large file systems with access to up to 14 TB of storage. Each of today's 8-node, 14-way SMP *AlphaServer* configurations supports 28 GB of memory. And next-generation 32-way *AlphaServer* systems with 128 GB memory will extend this scalability exponentially.

"*TruCluster Server* software on *AlphaServer* systems gives

us the flexibility we need by allowing us to increase compute power or storage simply, by adding more nodes or more storage to the cluster," says Morris. "And there's no need to shut down the cluster or interrupt service to do it."

A single system image for dramatically simplified management

Compaq's *TruCluster Server* software on *AlphaServer* platform not only delivers industry-leading availability, performance and scalability — it significantly reduces the complexity and effort involved in managing high availability computing solutions.

That's because *TruCluster Server V5* is a true single image system. Members of the cluster operate as a single system, transparently sharing resources, including data storage and clusterwide file systems — yet, at the same time, individual nodes can boot or shut down independently, without disrupting the cluster.



▲ *StorageWorks RAID Array*

As a result, Genentech's IT staff need to perform management operations only once — not over and over again for each system in the cluster. And they can use the same tools they use to manage single systems. In fact, because the cluster operates as a single system, it is actually less work to manage than separate standalone systems.

A fully shared clustered file system (CFS) enables all systems in the cluster to share common root, /usr, and /var file systems.

"The single root means I only have to edit a file once for the entire cluster," says Morris.

Disks and files are equally accessible from all nodes, at speed, concurrently and without specific preparation,

management, performance and reliability enhancements offered with *Tru64 UNIX V5* and *TruCluster Server V5*. Adding a new node into the cluster while online takes less than 30 minutes and the additional management overhead is practically nil since *TruCluster Server V5* is a true single-system image environment, substantially reducing management complexity."

"In addition to being able to script the startup, shutdown and application health, CAA also makes it easy to link applications to other applications, as well as network services. This linking of applications had proven very powerful in our Web-based deployments where we often have a number of independent applications dependent on the availability of the Web server."

"Single system image is the key to improved manageability," says Morris.

"From any node, I can access all of my storage, configuration files and so on — without the overhead and headaches associated with NFS."

Operating system and application software is installed just once for the entire cluster. Because the cluster operates as a single domain, there's no need to manually configure security privileges or other attributes that must be kept in synch across all servers. And because the cluster appears as a single IP address on the network, clients don't have to connect to a specific node for services. If the node providing a service goes down, client requests are automatically re-directed to another system in the cluster.

regardless of whether the file is stored on a locally attached disk or on a shared storage bus. In the event of a failed storage or network controller, a node switches automatically to an alternate controller using multi-path I/O — or, if that path has also failed, it continues to work by remotely utilizing some other node's storage controller. This ability to share and manage storage cluster-wide can reduce overall storage costs by 50% to 90%.

"This stuff is seriously cool," says Morris. "We are extremely impressed with the manage-

Managing anytime, anywhere

Beyond single system image, management is further simplified by intelligent, on-board event management that helps monitor, interpret and respond to events quickly, before they can impact system performance.

"The new Event Management System (EVM) is a significant addition to our ability to monitor and manage a large system," say Morris. "It provides tools for filtering, sorting, displaying, and in extreme cases, broad distribution of event information about the operation of the system, storage and applications. Because EVM can incorporate information from any text or syslog-based logging, it provides an integrated view of the entire system."

A new Cluster Application Availability (CAA) facility provides the same functionality as failover scripts, but goes much further, says Morris.

Online operations

Online storage management, tuning and configuration capabilities make it easy to make changes across the cluster — without interrupting operation. Genentech plans to use these facilities of *Tru64 UNIX V5* and *TruCluster Server V5* to support their Intranet.

"*TruCluster Server V5* makes it very easy to run instances of Netscape Enterprise Server, now called iPlanet Web Server, on multiple nodes of the cluster at the same time," says Morris. "The dynamic access capability inherent with *TruCluster Server V5* transparently directs Web clients to alternate nodes, providing the functionality of a network load-sharing device without the complexity or data synchronization such approaches used to require."



“We’ve long relied on Compaq solutions,” says Morris. “Compaq Services provides us with fast, knowledgeable support and maintenance of our systems. The *AlphaServer* is a very hot box. And *Tru64 UNIX* and *TruCluster* software give us a lot of application flexibility while ensuring the kind of high availability we need. Now, with V5, we get a significant improvement in manageability.”

What’s next at Genentech?

“Our next project is the migration of our Oracle databases, used for everything from financials to sales information to laboratory data, to a highly available clustered environment built around *AlphaServer*, *TruCluster Server V5* and *StorageWorks Fibre Channel* arrays,” says Morris.

Morris expects this cluster to be the core of Genentech’s three-tier database environment and to deliver substantial improvements in performance and availability. “We believe we will also achieve a substantial reduction in overhead by migrating many of our Oracle instances to a single, highly available back-end,” he says.

For more information

To learn more about *Tru64 UNIX* on *AlphaServer* solutions or locate the name of your nearest Compaq sales representative or authorized business partner, visit our website at www.tru64unix.compaq.com

COMPAQ, NonStop, AlphaServer, StorageWorks, Tru64, TruCluster Registered in U.S. Patent and Trademark Office.

All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Compaq Computer Corporation shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.

12A8-0500A-WWEN REL 000 04 27 5,0 Copyright 2000 Compaq Computer Corporation. All Rights Reserved.

COMPAQ

