The University of Florida speeds up memory intensive gene research with Dell HPC solution

- Blade Solutions
- High-Performance Computing

"By publishing papers, our scientists use the data from the Dell VSMP cluster to generate the next round of proposals that will attract funding."

Bill Farmerie,
Associate Director, Interdisciplinary Center for Biotechnology Research, University of Florida

Customer Profile

Company:	University of Florida College of Medicine
Industry:	Education
Country:	United States
Students:	1,642
Faculty:	932
Web:	www.med.ufl.edu

Business Need

The University of Florida Interdisciplinary Center for Biotechnology Research (ICBR) needed to speed up its HPC processing power and increase memory to keep up with the pace of gene sequencing, which produces a million times more data than it did two decades ago. Researchers had applications which took too long to run or could not run at all due to memory limits on the university's existing HPC clusters.

Solution

The university deployed 16 Dell™ PowerEdge™ M610 blade servers configured with Intel® Xeon® processors 5500 series, ScaleMP software solution and a quad data rate (QDR) InfiniBand interconnect to create a high performance virtual symmetric multiprocessing (VSMP)-based platform for genomics sequencing. The solution cluster has run previously memory-constrained jobs up to 160 times faster than its previous HPC solution.

Benefits

- **Eliminated disk swapping, allowing memory hungry jobs to run up to 160x faster (1½ hours vs. 10 days)**
- **40% performance improvement with Intel Xeon 5500 series and 5600 series processors compared with previous generation**
- **3-hour setup of blade enclosure, including blade servers**
- **1-day installation of ScaleMP software**
- **Lower power and cooling costs**
- **Support for more research projects and more grants**

The whirlwind speed of progress in the computer industry happens at an exponential rate that can be predicted. According to Moore's Law, the processing speed, memory capacity and even the number of pixels in digital cameras doubles every two years. But there is a branch of technology that's evolving even faster: gene sequencing, the ordering of nucleotides that make up a strand of DNA in an organism. Gene sequencing is the basis of the whole group of life sciences that study the genetic makeup of humans and other organisms in order to extend life.

"We were hitting the memory limits on individual nodes, which meant that jobs took longer and sometimes failed. One assembly application had taken 10 days on our old cluster, and it took only an hour and a half to complete on the Dell VSMP cluster."

Aaron Gardner, Cyberinfrastructure Director, Interdisciplinary Center for Biotechnology Research, University of Florida

Compared with capillary-based sequencing technology of just a few years ago, today's next-generation sequencing is able to produce a million times more data, which drives up the demands for computation and storage.

"We've gone from first-generation DNA sequencing instruments in the 1990s that analyzed 384 sequences at one time to instruments deriving 400 million sequences in parallel," says Bill Farmerie, associate director of the Interdisciplinary Center for Biotechnology Research (ICBR) at the University of Florida. "And as the volume of data is growing exponentially, that cost of data production has come down by a factor of 100,000."

Speeding up the pace of biotechnology research

The University of Florida is working to satisfy the demand for faster computers in its ICBR. There, scientists are working with Dell and Intel technology to construct the next generation of high performance computing (HPC) clusters that can keep up with the computational needs of the gene sequencing industry.

"We need to attack larger problems, larger genomes, larger samples and just get larger views of the systems we are studying," says Aaron Gardner, cyberinfrastructure director, Interdisciplinary Center for Biotechnology Research, University of Florida.

As the query sets and databases the queries are run against grew over time, the amount of memory that was available on a computer became of paramount importance. For best performance, it was necessary to cache numerous databases in memory and parallelize the algorithms being used so that they could all share memory between the nodes. The concept of symmetric multiprocessing (SMP) in the HPC cluster evolved to become virtual symmetric multiprocessing (VSMP), in which multiple physical systems appear to function as a single logical system.

Achieving very large shared memory

"We found that traditional HPC on a cluster wasn't working because we had hard limits on how much memory was available on a single node, and often the software was ill equipped to be able to distribute these databases across all the nodes in a cluster," says Gardner. "The VSMP system allows us to have a very

Technology at Work

Hardware

- Dell™ PowerConnect™ 6248, 6224, 6220, 5224 Gigabit Ethernet switches
- Dell PowerConnect 8024 10 Gigabit Ethernet switches
- Dell PowerEdge™ M610 blade servers with Intel® Xeon® processors 5560
- Dell PowerEdge M1000e modular blade enclosure
- Mellanox M3601Q QDR InfiniBand blade switch

Software

- Red Hat Enterprise Linux 5 with ScaleMP optimized kernel for RHEL
- ScaleMP vSMP Foundation for SMP

large shared memory space where we can cache in memory all of the sequence databases and all of the queries that we are searching against. This makes them accessible to all the processors at the same time with minimal latency."

To build the cluster, the ICBR populated a Dell PowerEdge M1000e modular blade enclosure with 16 Dell PowerEdge M610 blade servers with Intel Xeon processors 5560 and DDR3 memory, which provides approximately one terabyte of shared memory. A quad data rate (QDR), 40 Gb per second Mellanox M3601Q InfiniBand blade switch sourced through Dell busses all the memory and CPU calls between blades.

"We specifically waited for the Intel Xeon processors 5500 series to be developed because of the Intel QuickPath technology which enables all the cores on the individual CPUs, as well as the adjacent sockets within a system, to more quickly route data between their caches," says Gardner. "Using DDR3 memory with Intel Xeon processors 5500 series is a good match because the processor has higher available bandwidth and memory interface. When we paired the Intel processor with DDR3 with QDR InfiniBand, we were able to minimize latency and improve throughput in the VSMP system for memory performance. The Intel Xeon processors 5500 series alone give us a raw performance improvement of 40 percent up from the previous generation of Intel processors, so we're building our system on a much faster processor."

ICBR chose the Dell PowerEdge M1000e blade chassis for the VSMP system for multiple reasons. "It was the only system that we considered that could get us the buffered DDR3 DIMMs that we needed within our time constraints," says Gardner. "Of the systems we considered, it was also the only one available with QDR InfiniBand, and it facilitates the InfiniBand interconnect between the nodes using the backplane, so there are no cables involved. That increases the reliability and uptime of the VSMP system. So the Dell system was the most complete system, feature wise, for deploying the VSMP solution, as opposed to the others we considered."

Reducing management overhead

The dual Dell Chassis Management Controllers (CMC) in the PowerEdge M1000e modular blade enclosure provide redundant, secure access paths for administrators to manage the blades from a single console as a single system image. Integrated Dell Remote Access Cards for all the blades and enclosure enable remote management, which, along with reduced complexity on the management end, helped to give Gardner's team more time to work with researchers on how to best utilize the resources.

"Another factor that we like is the power footprint," says Gardner. "The Dell PowerEdge blade system has only six power supplies, three of which are required to run the system, and those are higher efficiency power supplies. It helps us to save 6U-10U of rack space and also save on the limited resources we have in our server room for power and cooling versus having power supplies in each discrete system."

The VSMP technology itself is provided by Dell Business Partner ScaleMP. With ScaleMP vSMP Foundation for SMP software, multiple physical systems appear to function as a single logical system. The innovative ScaleMP Versatile SMP (vSMP) architecture aggregates multiple x86 systems into a single virtual x86 system, delivering an industry-standard, high-end SMP computer. ScaleMP uses software to replace custom hardware and components to offer a new, revolutionary computing paradigm.

Fast setup and deployment

Once ICBR received the Dell enclosure and blades, it took Gardner and his team about three hours to get it racked and powered up and do the diagnostics. "We were taking our time," says Gardner. "We could have done it faster."

Deploying the VSMP software took about one day with a ScaleMP representative on site facilitating the VSMP technology. "We were able to accomplish that because the Dell hardware functioned without a hitch," says Gardner. "And also because we had already created a standardized hardened production image that we were able to deploy on the system. We've purchased a lot of hardware

"The Dell system was the most complete system, feature wise, for deploying the VSMP solution, as opposed to the others we considered."

Aaron Gardner, Cyberinfrastructure Director, Interdisciplinary Center for Biotechnology Research, University of Florida

from Dell in the past, so it was very easy for us to work within the existing relationship and arrive at the VSMP solution quickly."

Up to 160x faster results

Prior to the VSMP solution, there were several applications that the university was running on a standard SMP cluster of x86 machines with unsatisfactory results. "We were hitting the memory limits on individual nodes, which meant that the jobs took longer and sometimes just failed," says Gardner. "So having this larger memory system has enabled us both to get jobs done and to see them through to completion. We've seen some substantial performance improvements because we're able to run all of the data in memory, without going to disk. For example, one assembly application had taken 10 days on our old cluster, and it took only an hour and a half to complete on the Dell VSMP cluster."

In addition, the university is able to allow researchers to run these jobs in an interactive, real-time manner, rather than waiting in a queue. This enables them to experiment more, and even develop new analysis methods with the system. "This really helps to achieve a better answer in terms of the analysis we're performing," says Gardner. "For instance, a DeNova assembly application was swapping to disk because all of the sequences and alignments could not fit into memory. Before we had the VSMP system, we would have had to resort to discarding some sequences, or assembling several smaller assemblies together. These approaches can sometimes produce inferior or misleading results, and cause you to lose consistency and depth with the assembly statistics the software captures. By being able to successfully assembly all of a project's data at once and get the result back quickly, we are empowered to iterate parameters and adapt our analysis methods in near real time. This is preferable to waiting a week and being forced to work with whatever we get due to the time constraints involved in rerunning those types of jobs."

Completing the circle

In addition to the Dell VSMP solution, ICBR has been purchasing Dell PowerConnect 6248, 6224, 6220, 5224 Gigabit Ethernet switches for networking infrastructure. "We like the stacking capabilities of the Dell switches," says Gardner. "We've also purchased some Dell PowerConnect 8024 10 Gigabit Ethernet switches as a front-end interconnect to replace our existing Gigabit infrastructure. With recent 10GbE hardware we are starting to see the performance improve to an acceptable level and we can run almost any protocol over 10GbE. We can also pull a larger portion of our computing staff into supporting research computing because the networking and storage protocols and paradigms with 10GbE are familiar to them."

As ICBR provides its researchers with the faster processing power they demand, the science of gene sequence will speed up and produce faster results for the life sciences. The immediate result will be more research projects and more grants to fund them.

"It's all in the papers," says Farmerie. "By publishing papers, our scientists use the data from the Dell VSMP cluster to generate the next round of proposals that will attract funding. So there's the cycle that has to be completed each time in order to drive the process of science further down the road."

For more information go to:
www.DellHPCSolutions.com

View all Dell case studies at: dell.com/casestudies