

Cisco HPC Network Aids in Complex Bio-Molecular Simulations at the University of Florida

The University of Florida uses high-performance computing to simulate protein folding and help in the fight against disease.

Challenge: The Protein Folding Problem

Just like a road map, there are many ways to fold a protein molecule but only one is right. Misfold a map and the only penalty is inconvenience; but misfold a protein and the penalty can be a bad

EXECUTIVE SUMMARY
<p>UNIVERSITY OF FLORIDA</p> <ul style="list-style-type: none"> Quantum Theory Project and Department of Chemistry Gainesville, Florida, United States <p>RESEARCH CHALLENGE</p> <ul style="list-style-type: none"> Understanding the origins of disease by simulating the motion of molecules and folding of proteins <p>SOLUTION</p> <ul style="list-style-type: none"> Low-latency, high-bandwidth InfiniBand solution for interconnecting the 200 node (800 processor) HPC cluster High-speed Ethernet solution for management connectivity and access to the research network High speed Fibre Chanel Block storage connectivity to 45-TB storage system <p>BUSINESS RESULTS</p> <ul style="list-style-type: none"> Researchers can now tackle fundamental problems through "brute force" simulation approach that would have previously taken years High-speed infrastructure serves many different research projects

disease. How does a protein know the shape into which it is supposed to fold? High-performance computing can help answer this question. Figure 1 shows an energy landscape and a cartoon of an unfolded, transitional, and folded protein. Low free energy is good. Laboratory experiments can probe around only the unfolded and folded regions of the energy curve. Computer experiments can probe the whole thing.

Seonah Kim, a graduate student in chemistry at the University of Florida, and Professor Adrian Roitberg are doing just that on the UF High Performance Computing (HPC) Cluster. The cluster depends on the high performance and reliability of the Cisco® InfiniBand fabric that connects the AMD Opteron based Rackable servers and storage subsystem. Kim has run more than 45 days on 100 processors and isn't done yet.

The simulation uses the highly parallelized Assisted Model Building with Energy Refinement (AMBER) package of molecular simulation programs. Why so long to study just two proteins? For one thing, biology involves a lot of water. Figure 2 shows a visualization snapshot of the computed system. The pinkish cloud is 7000 water molecules (21,000 atoms) surrounding a 14-residue peptide molecule (the bluish "worm" in the middle). The AMBER simulation works by calculating the motions of all these molecules. They bend, rotate, and move through space, avoiding or bouncing off one another. The simulation divides time into little steps and uses Newton's laws of physics to calculate the motion of the thousands of atoms at each step.

If that weren't enough, there is the problem of "trapping." The energy landscape cartoon of Figure 1 is too simple. Real energies depend on molecular configuration in much more complicated ways, with many shallow energy valleys. The computation can get stuck in them, so multiple replicas are run at a span of temperatures, and every so often the replicas are forced to jump to a new temperature. See Figure 3. This is very computationally intensive.

The University of Florida Research Team

Seonah Kim and Adrian Roitberg are members of the UF Quantum Theory Project, one of the largest and best known physics and chemistry interdisciplinary institutes in the world devoted to theory and computation of molecular and materials properties. More information is available at <http://www.qtp.ufl.edu>.

The High-Performance Computing Initiative at UF is an innovative approach to such needs. The design is a computing grid, linking specialized research computing clusters to a central parallel cluster over a dedicated high-speed network. Funding from the National Science Foundation and a cooperative agreement with Cisco provided the routers and switches for that grid.

Figure 1. Energy landscape of unfolded, transitional, and folded proteins.




Figure 2. Visualization of 7000 water molecules surrounding 14-residue peptide molecule


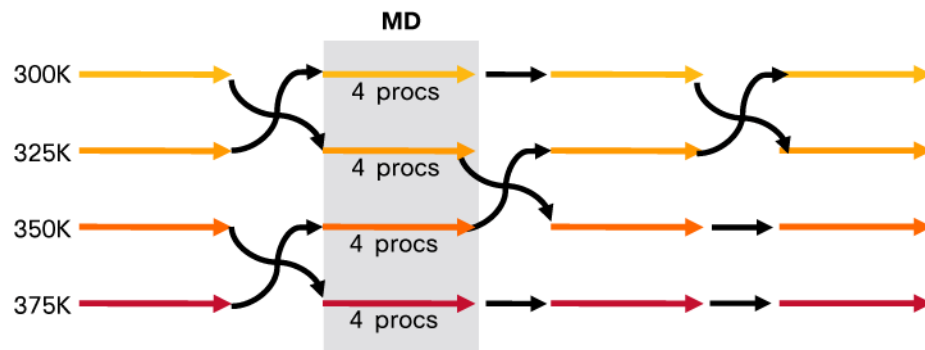


Figure 3. Molecular Exchange Simulation showing temperature jumps



- Molecular Dynamics (MD)
- 26 replicas, 4 processes/replica = 104 processes
- 650 sec/10 ps MD → 18 hr/1 ns MD → 1800 hr/100 ns MD = 75 days

Network Solution

UF operates an extensive network infrastructure, built using Cisco equipment and software, linking the various research facilities inside and outside the university. The production core network links the UF campus with the Internet and Internet2 networks, in addition to the UF Research Network. The UF Research Network is linked by 10 Gigabit Ethernet to the Florida Lambda Rail and in turn to the Ultralight network and National Lambda Rail networks.

In 2003, the University of Florida embarked on a campuswide grid strategy to support data-intensive research across multiple and diverse disciplines. The project was to enable researchers to achieve the full benefits of this grid strategy by adding dedicated high-performance interconnections and true data-intensive storage facilities. The major campus research computing facilities were to be linked by a 10-Gbps network and supported by 45 terabytes of high-performance storage systems.

The UF campus-grid and HPC infrastructure was intended to explore data-intensive, high-performance applications in six distinct research disciplines:

- High-energy physics
- Chemical physics and materials science
- Coastal and estuarine modeling
- Medical physics
- Computational biology
- Computer science and engineering

Technical Implementation

The University of Florida required a network infrastructure capable of broad and varied requirements:

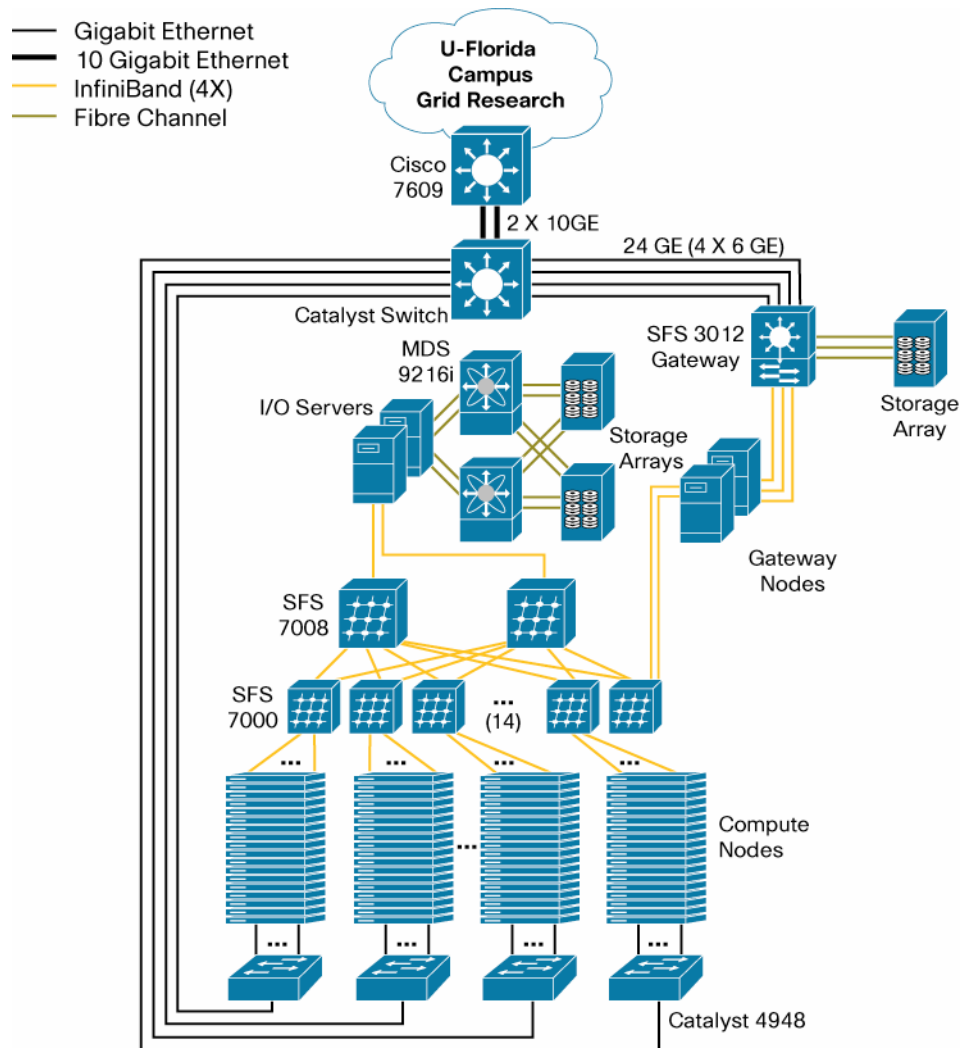
- **Scale:** Ability to scale an HPC cluster to 400 multiprocessor nodes and beyond
- **Broad application support:** Ability to support a wide range of applications from tightly coupled to massively parallel and parametric (such as Monte Carlo simulation)

- **High-performance MPI network:** A high-bandwidth, low-latency Inter-Process Communications or Message Passing Interface (MPI) network for peak performance of tightly coupled applications
- **High-performance storage network:** Ability to support high-bandwidth connectivity to a parallel file system and a back end Fibre Channel attached block storage

UF settled upon the following HPC cluster networking configuration.

- For the MPI network, UF based its solution on Cisco SFS 7000 Series InfiniBand Server Switches, using two 24-port switches per rack. Each rack consisted of 32 Rackable Systems compute nodes. Sixteen ports from each switch were connected to servers and the remaining eight ports connected as uplinks to two core 96-port Cisco SFS 7008 InfiniBand Server Fabric Switches (see Figure 3). The InfiniBand network is also used for high-bandwidth connectivity between the compute nodes and I/O nodes using Sockets Direct Protocol.
- Cisco Catalyst® 4948 Switches for the cluster management network and a Cisco Catalyst 4948 10 Gigabit Ethernet switch for outside cluster access and connectivity to the campus grid network and worldwide research networks.
- Cisco MDS 9216i Multilayer Fabric Switches for high performance and scalable block storage fanout from the six I/O nodes to the Xyratex storage arrays.
- A Cisco SFS 3012 Multifabric Server Switch for multiprotocol gateway capability between InfiniBand, Fibre Channel, and Gigabit Ethernet. The switch provides SCSI RDMA Protocol support from the InfiniBand-connected I/O nodes to the Fibre Channel-connected storage arrays, making the I/O nodes appear to have a direct SCSI block mode attachment to the storage arrays. Through the switch's Gigabit Ethernet gateway capability, UF is also able to access additional storage and datasets through the campus grid network.

Figure 4. UF HPC Cluster Architecture featuring networks for Ethernet, InfiniBand, and Fibre Channel



PRODUCT LIST

Ethernet Switching

- Cisco Catalyst 7609 Switch
- Cisco Catalyst 4948 and 4948-10GE Switches

InfiniBand Switching

- Cisco SFS 7000 Server Fabric Switch
- Cisco SFS 7008 Server Fabric Switch
- Cisco SFS 3012 Multifabric Server Switch

Storage Networking

- Cisco MDS 9216i Multilayer Fabric Switch

For More Information

To learn more about UF's HPC initiative, go to:

<http://www.hpc.ufl.edu>

To find out more about Cisco High Performance Computing, go to: <http://www.cisco.com/go/hpc>

**Americas Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)