

Cisco HPC Network Aids in Dark Matter Search at University of Florida

The University of Florida uses high-performance computing clusters in a search for the universe’s missing dark matter.

EXECUTIVE SUMMARY
<p>UNIVERSITY OF FLORIDA</p> <ul style="list-style-type: none"> • Department of Physics • Gainesville, Florida, United States
<p>BUSINESS CHALLENGE</p> <ul style="list-style-type: none"> • Understanding the early structure of the universe.
<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 Channel Block storage connectivity to 45-TB storage system
<p>BUSINESS RESULTS</p> <ul style="list-style-type: none"> • Researchers can now tackle fundamental problems with “brute force” approaches that would have previously taken years • High-speed infrastructure serves many different research projects

Challenge: Searching for Dark Matter

In the last few years, astrophysicists and astronomers have found strong evidence that most of the mass and energy in the universe is “dark”—no light or radio waves or other electromagnetic radiation are emitted. Normal observations (telescopes for example) miss this mass and energy. Nevertheless “dark matter” influences observable objects by its gravitation. But what is it? The best candidate is some kind of as-yet unobserved elementary particle, known as WIMP, for Weakly Interacting Massive Particle.

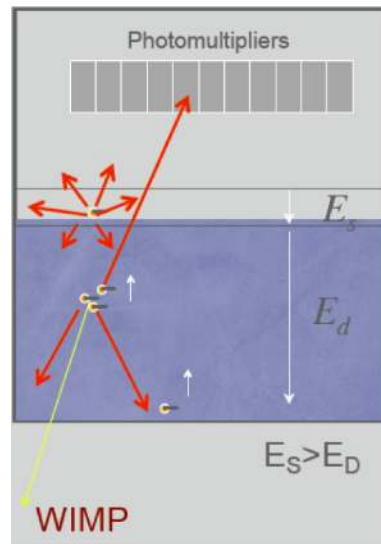
Building detectors for these particles faces huge problems. WIMPs would barely interact with normal matter so they would have very weak signals. Worse, all normal matter has many radiation sources that can give false signals or “background.” The dark matter signal is the needle in a haystack of background signals.

The solution? Model the background on a high-performance computing (HPC) cluster.

Professor Tarek Saab is working on this problem at the University of Florida (UF). One detector uses highly purified liquid xenon. Occasionally a WIMP will hit a xenon atom and cause it to emit light. Detected by phototubes, the light would reveal the WIMP’s interaction point. Electrons would also come off and then emit light, so they would also help trace the WIMP’s location.

See Figure 1.

Figure 1. Detecting Light Signals in the Xenon Experiment



Simulating the background “haystack” is statistical, which means lots of computation. Saab is using the UF High Performance Computing Cluster to achieve levels of insight and completeness not previously available. His calculations depend on the high reliability of the Cisco® InfiniBand fabric that connects the AMD Opteron–based Rackable servers and storage subsystem. Thousands of runs proceed for days. The simulations use models of background radiation sources from the various detector components in their proper geometry. From physical principles, the tracks of all these can be calculated. Hundreds of thousands of tracks are done. The results are shown in Figure 2. The simulation matches the measured background. If a signal doesn’t match the simulation, it should be studied carefully. It might be the elusive WIMP.

Monte Carlo Simulation

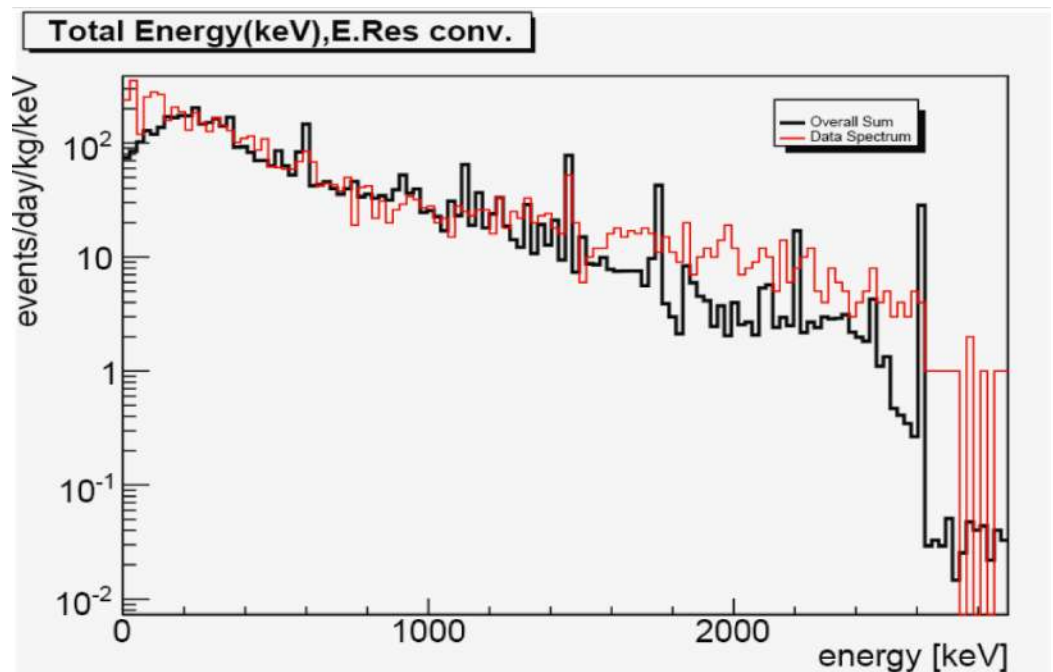
The experiment uses a Monte Carlo simulation on the UF HPC Cluster. A Monte Carlo simulation randomly generates values for uncertain variables over and over to simulate a model. The amount of time required to run a simulation is dependent upon which decay is simulated and which geometry is used. The number of events per simulation in this experiment ranges between 10 million and 10 billion, depending upon the decay. To get good statistics, it takes one to three days for all the simulations to complete.

Research at the University of Florida

The High-Performance Computing Initiative at UF is an innovative approach to research. The design is a computing grid that links 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.

Professor Tarek Saab is a member of the experimental particle astrophysics effort in the UF Department of Physics. For more information on Saab's work, see <http://www.phys.ufl.edu/~tsaab>.

Figure 2. Hamatsu PMT Simulation. The difference between the Overall Sum and the Data Spectrum represented by the red and black curves could reveal the elusive WIMP.



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 backend Fibre Channel attached block storage

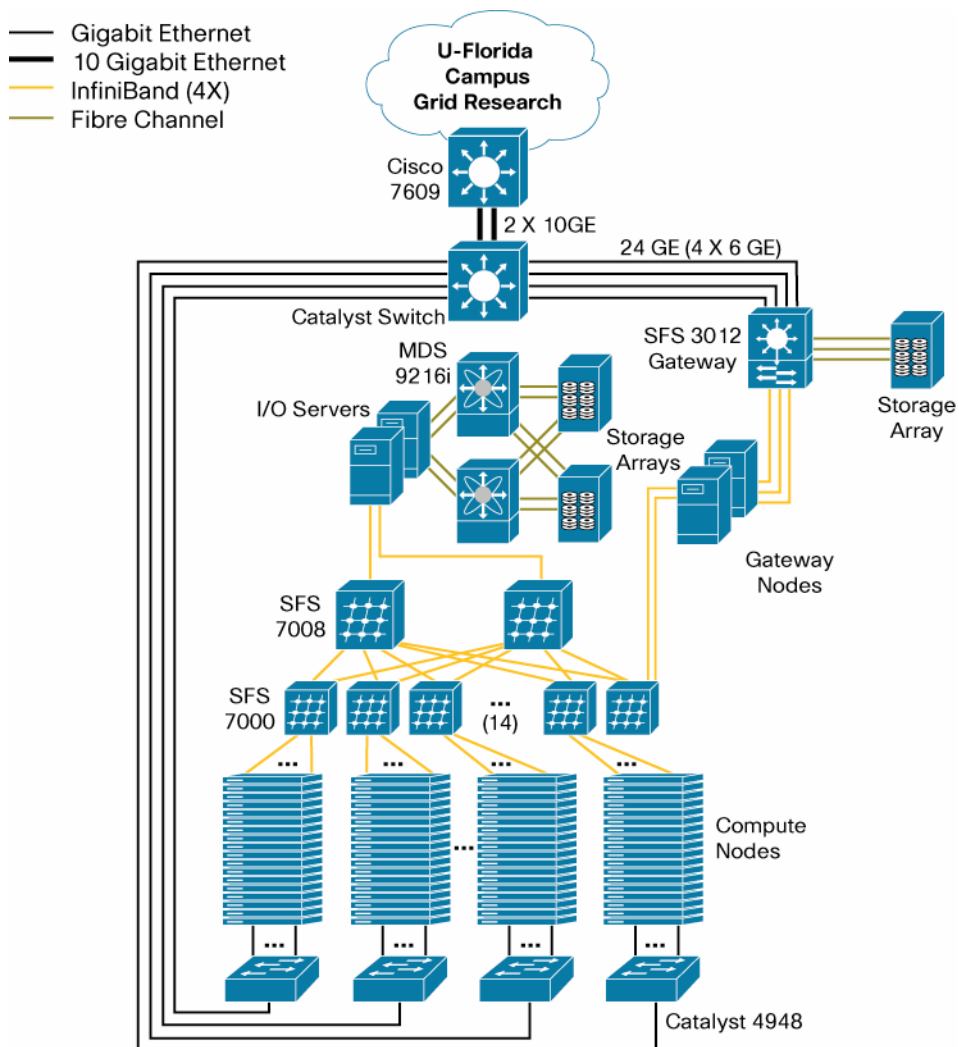
UF settled upon the following HPC cluster networking configuration:

PRODUCT LIST
<p>Ethernet Switching</p> <ul style="list-style-type: none"> • Cisco Catalyst 7609 Switch • Cisco Catalyst 4948 and 4948-10GE Switches
<p>InfiniBand Switching</p> <ul style="list-style-type: none"> • Cisco SFS 7000 Server Fabric Switch • Cisco SFS 7008 Server Fabric Switch • Cisco SFS 3012 Multifabric Server Switch
<p>Storage Networking</p> <ul style="list-style-type: none"> • Cisco MDS 9216i Multilayer Fabric Switch

- 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 3. UF HPC Cluster Architecture featuring networks for Ethernet, InfiniBand, and Fibre Channel



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)