

# University Pioneers Grid Computing and High-Speed Networking

Multidisciplinary teams at the University of Florida gain 100x increase in compute and network throughput for collaborative research.

## EXECUTIVE SUMMARY

### UNIVERSITY OF FLORIDA, HIGH PERFORMANCE COMPUTING CENTER AND CAMPUS GRID PROJECT

- **Industry:** Higher Education
- **Location:** Gainesville, Florida

### BUSINESS CHALLENGE

- Facilitate faster computations and more conclusive results from multidisciplinary research
- Collaborate in design and implementation to create a comprehensive leading-edge solution for a reasonable cost
- Provide leading facilities that strengthen the university's competitive position within global research and education communities

### NETWORK SOLUTION

- Phase 1: compute grid with Cisco Server Fabric Switches for high-speed interconnections
- Phase 2: second, larger grid and the introduction of high-speed campus network and links to off-campus research networks
- Continued collaborative research by University of Florida technology leaders and Cisco engineering to investigate and develop new solutions for HPC environments

### BUSINESS RESULTS

- Shortened simulation times: 100x increase in on-campus compute power and network throughput
- Broadened base of researchers who can take advantage of shared HPC grid and networks
- Centralized, consolidated compute resources managed by the HPC Center
- Industry-leading facilities, strengthening participation in the Open Science Grid (OSG)

## Business Challenge

The research and education (R&E) community not only relies on high-performance computing (HPC) and networking, but also contributes to the advancement of these required foundational technologies. Research teams have traditionally designed, deployed, and operated their own infrastructures. However, these silos of resources force researchers and educators to divert attention from science to computing, and present escalating challenges in terms of support and cost as system and network complexity increases. As a result, university information technology (IT) teams have evolved to fill a crucial role on today's campuses.

"Efforts to coordinate HPC resources at the University of Florida began many years ago, and evolved from an HPC committee into the HPC Center," says Marc Hoit, CIO for the University of Florida. "It made a lot of sense for IT to coordinate resources across teams and departments. We could avoid redundant efforts, identify the common requirements, and provide an infrastructure that could greatly exceed the capabilities of any isolated cluster or network."

As the HPC Center was being established to overcome organizational barriers on campus, the physics department was driving demand for shared computing grids and high-speed interconnections. "Simulation is a vital part of scientific research," says Paul Avery, a scientist in the physics department at the University of Florida. "Within the physics community, we have been dependent on shared computing for quite some time, and now we are seeing many other disciplines – chemistry, biology and medicine, and earth sciences – that can use compute grids and high-speed networks to reduce simulation times from years down to weeks."

The grassroots efforts spearheaded by faculty came together, with support from IT and the administration, to overcome the challenges of moving to a next-generation model for campuswide grid computing. The specific requirements included compute and data movement solutions to support large-scale simulations. Additionally, campus networks and data movement solutions must also tie into national networks to facilitate collaborative research projects extending out to national and global R&E communities.

**"Our partnership with Cisco is essential.... Our on-campus capabilities have been increased by a couple orders of magnitude, and our collaborative findings will help improve future generations of commercially available solutions."**

– Erik Deumens, Director of the HPC Center, University of Florida

## Network Solution

The University of Florida Campus Grid computing initiative was carried out in phases:

- Phase 1 – General grid research within the physics department. The beginnings of the shared computing model involved cross-discipline funding – from physics faculty, the College of Liberal Arts and Sciences, and the CIO. A cluster supported by these sources and the International Virtual Data Grid Laboratory (iVDGL) project was deployed during Phase 1.
- Phase 2 – College of Engineering. The success of the first HPC cluster deployment inspired additional funding from the college of engineering faculty, administration, and the CIO, enabling a second larger cluster, with both InfiniBand and Gigabit Ethernet interconnects, as a shared resource for all of the engineering departments. This phase included a second faculty-driven initiative to augment grid computing solutions with a high-speed networking initiative, and included application to the National Science Foundation (NSF) for a grant to fund the prerequisite networking infrastructure.

Winning the NSF grant was vital for Phase 2, but the project's overall success required additional support. University of Florida, with its networking partner, Cisco®, has participated in a joint research project to invest in the design, deployment, and study of server fabric interconnections for the clusters, as well as the high-speed cluster-to-cluster connections and links to Florida LambdaRail (FLR) and National LambdaRail (NLR). With the long-standing relationship between Cisco and the university and the prior deployments of Cisco networking solutions, both on campus and within FLR and NLR infrastructures, the joint project successfully utilized the existing infrastructures to introduce a solution with unprecedented reliability, capacity, and performance.

“Our partnership with Cisco is essential,” says Erik Deumens, director of the HPC Center at the University of Florida. “By bringing together our experts in HPC and data management with Cisco’s best networking architects and engineers, we have accomplished something extraordinary and worth boasting about. Our on-campus capabilities have been increased by a couple orders of magnitude, and our collaborative findings will help improve future generations of commercially available solutions.”

“Lots of companies say that they want to partner with us, but what they mean is that they want to sell us products,” says Hoit. “With Cisco, we have a true partnership. We have both invested funds and resources, and both gained knowledge about new and novel research solutions and potential uses for the latest technologies within large-scale computing grids.”

The University of Florida Campus Grid computing initiative provides a platform from which researchers can participate in the Open Science Grid (OSG). Organized by the associated consortium of members who collectively deploy and own the shared infrastructure, the OSG spans approximately 40 sites in the United States, South America, and Asia. High-speed networks currently link more than 20,000 processors to provide a collaborative platform for multiple sciences. Members gain access to an infrastructure that could not be afforded by any one entity. The University of Florida, and Cisco as its partner, can use the OSG for both production and research projects in the sciences to test new grid technologies.

**“Our new grid infrastructure and connections with other high-speed R&E networks significantly strengthen our e-science position within the overall community. Together with Cisco, we have learned how to design and build better solutions.”**

– Marc Hoit, Chief Information Officer, University of Florida

## Business Results

With the completion of Phase 2, the university has given multidisciplinary teams access to much more compute power and the ability to move very large data sets among cluster resources both on campus and off. The new compute, storage, and networking capabilities translate into high-impact benefits for a broad range of multidisciplinary projects:

- **Ability to tackle large-scale problems** by providing researchers with access to hundreds of processors within the campus grid. Before this project, one group of researchers could not run their simulation without overwhelming the cluster. Since other teams were impacted, the group was not able to run their simulation. With the new facilities, the same group can run the simulation, and results are returned in minutes. In an hour, they are able to gain huge quantities of results.
- **Reduced processing times** for the most complex codes, saving project teams weeks and months, and speeding the time to results that affect many lives. For example, one project is helping in the fight against Mad Cow, Alzheimer's, and other diseases that result from misfolded proteins. To gain understanding and potentially discover breakthroughs, the team must simulate protein folds. A single protein-folding simulation can now be done in less than ten weeks, whereas the same simulation would have taken more than a year on the previously available compute platforms.
- **Access for more research teams**, not just the few that can afford the most elaborate HPC systems.
- **Continuing advancements in ease of use**, making the technology accessible to a growing number of scientists, in addition to the most computer-savvy that have pioneered clusters and grid computing.

“The value of the Campus Grid computing and networking initiatives directly affect researchers and the university as a whole,” says Hoit. “We have to compete for researchers and for funding. Our new grid infrastructure and connections with other high-speed R&E networks significantly strengthen our e-science position within the overall community. The infrastructure will enrich our research programs and also underscore our leadership in HPC and storage. Together with Cisco, we have learned how to design and build better solutions. We know where we can save money and where we need to spend money to remain a leader.”

The high-speed grids are changing the way that researchers work. In the past, data movement was a bottleneck. Scientists routinely simplified the problem until the required data set was manageable. Today, data sets can easily grow into hundreds of terabytes or even petabytes. The on-campus grids and access to the larger OSG and other national computing resources over high-speed interconnects make it possible to tackle problems using data sets in their entirety, and to

share large data sets among collaborative teams that can apply multiple viewpoints and parallel approaches.

“Two-thirds of what we do is social – connecting with other researchers is incredibly important,” says Avery. “Networks enable communications – video conferencing and other instant forms of relaying information. The latest breakthroughs in our high-speed campus network mean that we can move data from one place to another – even the largest data sets. When you can move data, you can bring more expertise to bear to understand the data and solve the problem. We see it happening already in medicine – doctors can send a half-gigabyte medical image to another doctor across the country. By enabling this type of convenient collaboration for molecular research, studying our environment, or discovering the origins of matter, we have opened the door for revolutionary breakthroughs in numerous fields of science.”

**“Any collaborative field like physics moves forward only as fast as its cyber-infrastructures allow. Our work with Cisco has given us the networking capabilities that we need to empower very large teams...the progress that we have made will directly affect global economies in years to come.”**

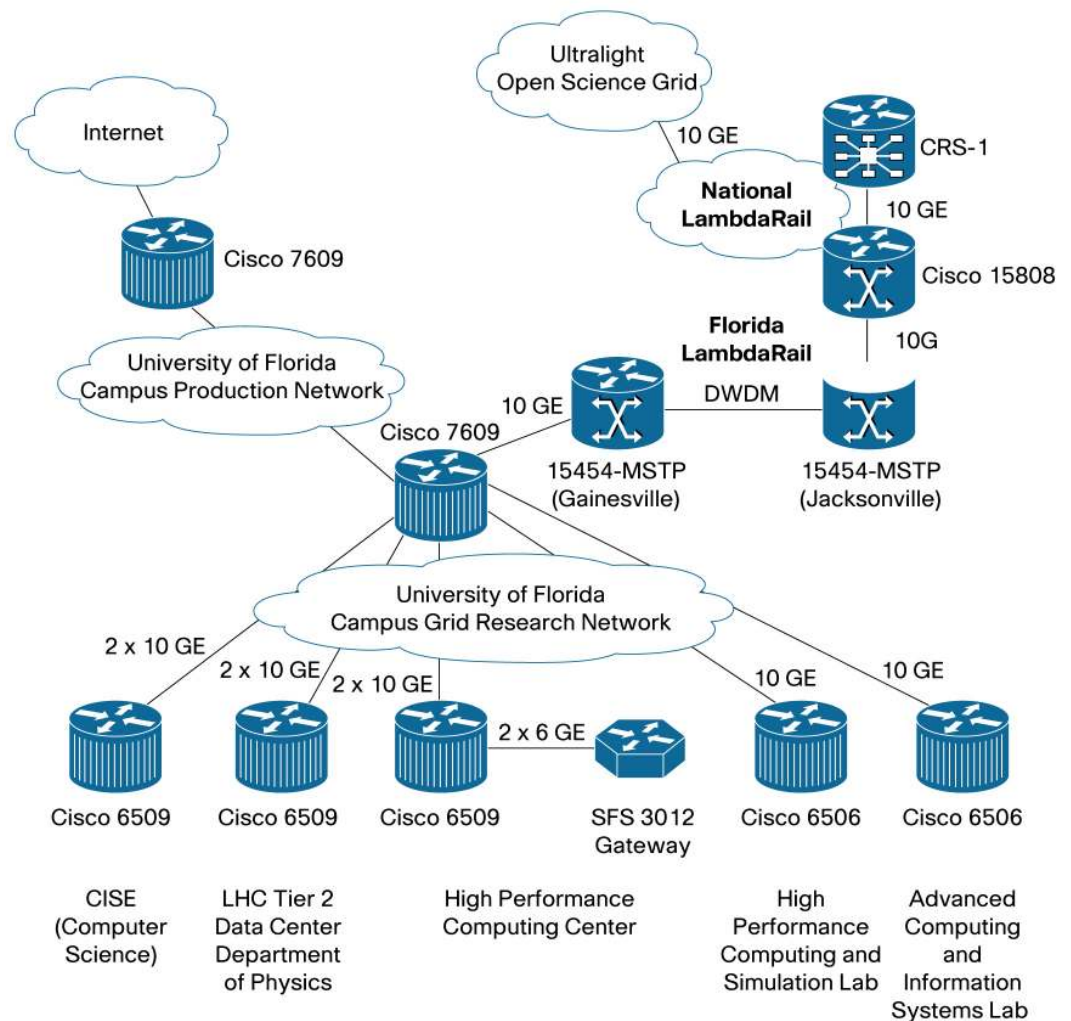
– Paul Avery, Physicist, Department of Physics, College of Liberal Arts and Sciences, University of Florida

### Next Steps

In Phase 3, grid computing and networking facilities will be made available to the health sciences center and life sciences department. These teams, while less computer-focused, will use a Web-like interface to deploy programs on the grid. By eliminating the need for researchers to understand intrinsic architectures and topologies, these innovations will simplify use for a broader number of researchers. The new interface will also allow changes to be made without affecting the user base. Development efforts are in progress, and this phase will be introduced to the campus in late 2007 and early 2008.

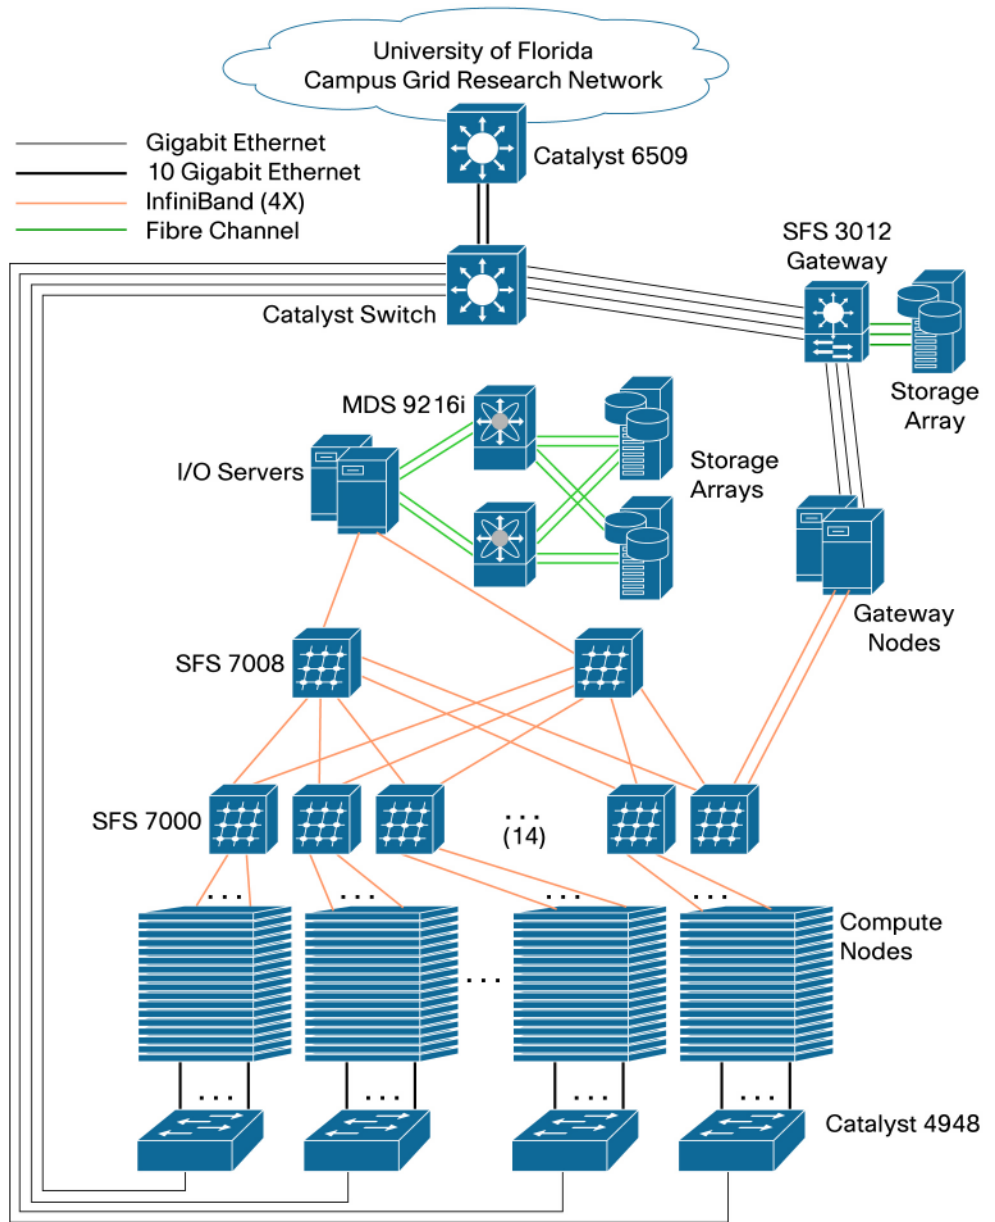
## Technical Implementation

**Figure 1.** Phase 1 and 2 Campus Grid



In Phase 2, the Campus Grid was expanded with a new cluster containing 200 nodes (see Figure 1). Each node includes two dual-core AMD processors, for a total of 800 processors. The InfiniBand server fabric was built using 14 Cisco SFS 7000 Series InfiniBand Server Switches (see Figure 2). A core level of switching, using two Cisco SFS 7008 InfiniBand Server Switches, connects all processors to 42TB of storage.

**Figure 2.** High-speed interconnections



The cluster delivers 1.4 GB/sec of processor-to-storage throughput today, and upcoming enhancements are expected to raise it to 2.5 GB/sec. This exceptional performance was achieved by the joint University-Cisco team, allowing a parallel file system to be deployed across an InfiniBand fabric of unprecedented scale and achieving performance that is close to the maximum throughput of the hardware components involved.

“The distributed data model is simple and cost effective,” says Avery. “Scientists have access to a file system of huge proportions that looks like it is on the desktop and operates at speeds equivalent to local storage. A single data set can reside in one location, and be streamed to any desktop. Data does not have to be copied or divided into subsets. And researchers can share a huge storage grid that could never be afforded by a single team.”

Historically, physics has always been a large-scale collaborative field. Even with much smaller data sets, the data management challenge dictated how research was done. The most challenging projects, with the largest number of researchers, had to rely on national laboratories to host data. Researchers had to travel to the laboratories to get full access to the data. The Internet and high-

### PRODUCT LIST

#### High-Performance Computing Solutions:

- Cisco SFS 7008 InfiniBand Server Switches
- Cisco SFS 3012 Multifabric Server Gateway
- Cisco Catalyst 6500 Series Switches
- Cisco MDS 9000 Storage Fabric Switches
- Cisco 7600 Series Routers

speed networks have made it possible to work remotely, but the performance compromises mean that remote researchers must work with data subsets and deal with connection latencies.

Data movement requirements have driven higher-speed networks, storage, and CPU farms. The University of Florida Campus Grid comes at a time

when the scale of projects continues to grow at a staggering pace. “Any collaborative field like physics moves forward only as fast as its cyber-infrastructure allow,” says Avery. “Our work with Cisco has given us the networking capabilities that we need to empower very large teams. Today that means hundreds or even thousands of researchers focused on a single experiment and the resulting data set. Research greatly affects many industries, and the progress that we have made will directly affect global economies in years to come.”

### For More Information

To find out more about the Cisco high-performance computing solutions, 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](http://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)