Meeting the Challenge of High-Energy Physics

How the Ultralight Consortium Is Finding Answers to the Universe's Oldest Questions

THE HIGH-ENERGY PHYSICS COMMUNITY IS CONDUCTING A new round of experiments to probe the fundamental nature of matter and space-time and to understand the composition and early history of the universe. These experiments face unprecedented engineering challenges due to the volume and complexity of their data and the need for collaboration among scientists around the world.

The massive, globally distributed datasets that will be acquired by these experiments are expected to grow to the 100-petabyte level by 2010, and will require data throughputs on the order of gigabits-per-second between sites around the globe. Although grid-based infrastructures developed by collaborations in Europe and the United States have provided massive computing and storage resources, they remain limited by their treatment of the network as an external, passive and largely unmanaged resource.

To overcome this limitation, major high-energy physics centers in the United States have formed the UltraLight consortium. The project, funded by the National Science Foundation, is based on a partnership that includes the California Institute of Technology, University of Florida, Florida International University, University of Michigan, Stanford Linear Accelerator Center, Fermi National Accelerator Laboratory, Brookhaven National Laboratory and CERN (in Geneva, the world's largest particle physics laboratory). Industrial partners recognizing the potential benefits for future network developments and commercial applications have also joined the project. Cisco's Academic Research and Technology Initiatives (ARTI) division and Level(3) are integral to the project—their participation is materialized through concrete technical and strong financial support. Specific networking resources presently being made available to UltraLight are shown in Figure 1. They include the major facilities of LHCNet[1], transcontinental 10 gbps wavelengths from National LambdaRail[2] and UCAID[3], and partnerships with StarLight[4]. Transcontinental and intercontinental wavelengths in our partner projects TransLight[5], Netherlight[6], UKlight[7], AMPATH[8] and CA*Net4[9] will be used for network experiments on a part-time or scheduled basis.

UltraLight will monitor, manage and optimize the use of the network in real-time using a distributed set of intelligent global services. These services will leverage grid middleware, network-aware applications and heuristic optimization algorithms to form an integrated system designed to meet the experiments' needs. The developments will be driven by and will progressively serve the needs for grid-based data analysis and production-processing. The UltraLight hybrid packet- and circuit-switched network infrastructure will employ ultrascale protocols and the dynamic building of optical paths to provide efficient fair-sharing on long-range networks up to 10 gbps. Recently, the UltraLight team used the new FAST TCP[10] algorithm to sustain a single TCP stream of 7.4 gbps over a span of two hours between Pasadena and Geneva—a distance of 11,000 kilometers (approximately 6,836 miles).

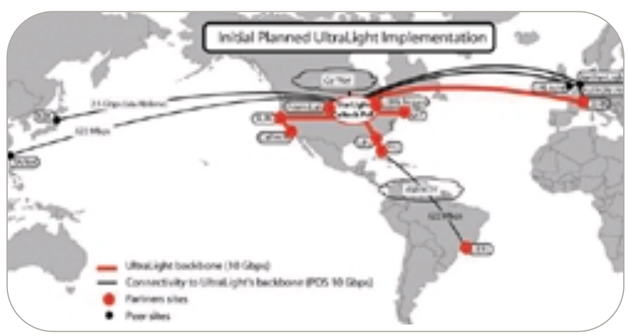

Figure 1: UltraLight Site Diagram

The UltraLight application-level services domain will provide interfaces and functionalities for the physics applications to effectively interact with networking, storage and computation resources as described in Figure 2. UltraLight will

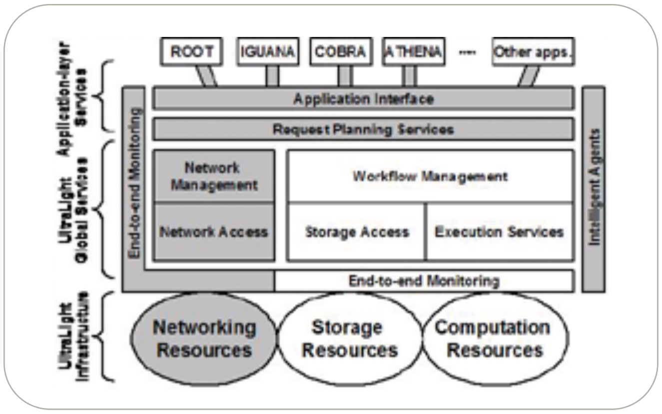

Figure 2: Context of UltraLight work. White regions represent areas which will be leveraged by UltraLight. Shaded regions represent areas that will be developed by UltraLight.

add a completely new dimension to the grid-based systems by interfacing applications to managed-networking services. It will extend the advanced planning and optimization behavior into the networking and data access layers, allowing a whole new class of advanced system behaviors and functionalities.

In addition, UltraLight plans to implement a set of global end-to-end managed services, building on the ongoing and rapidly advancing work on the Monitoring Agents in a Large Integrated Services Architecture system, or MonALISA[11]. The MonALISA system, shown in Figure 3, provides a distributed-service architecture used to collect complex monitoring information and processes it in a distributed-agent framework. The scalability of this agent-based framework is derived from the use of a multithreaded engine to host a variety of loosely coupled, self-describing dynamic services; and from the ability of each service to register itself and to be discovered and used by any other services or clients that require such information. The MonALISA system can be used to monitor and control different network devices, including photonic switches. Gathering the information collected from multiple points allows for the generatation of global views for connectivity and for the ability to spot problems and develop higher-level services for decisions. Mobile agents capable of providing optimized dynamic routing for distributed applications have recently been added to the system.

UltraLight's approach will allow applications and higher-level service layers to be made aware of advanced behaviors, as well as options available within the system, and to provide the required interfaces to grid middleware services. This enables a new class of proactive and reactive applications that can dynamically handle unexpected system behaviors, such as congestion or hardware failures, and allow for dynamic responses to changes in the system setup, such as when new network paths or modes become available. These new functions will enhance global system resilience to malfunction and allow optimization of resource usage, improving overall throughput and enabling effective implementation of policies.

A special and critical class of applications is those to be used for analysis of data from the Large Hadron Collider (LHC), the new particle accelerator being built at CERN. The Grid Analysis Environment (GAE)[12] being developed at Caltech and University of Florida will enable thousands of users to harness the full power of the grid. They'll be able to discover, analyze and collaborate on the petabytes of data generated by the LHC experiments through the use of grid portals based on grid (Web) services designed to hide much of the grid's infrastructure and resource complexity. Components of the GAE will interact with the MonALISA environment and replicate data, schedule jobs and find optimal network connections in an autonomous manner. The result will be a self-organizing grid with no single point of failure, in which thousands of users are able to get fair access to a limited set of distributed resources of the grid in a responsive manner.

continued on page 20

Figure 3: The MonALISA framework provides a distributed monitoring service system. Each MonALISA server acts as a dynamic service system and provides the functionality to be discovered and used by any other services or clients that require such information.

CPEC Revolutionizes Online Professional Development

continued from page 6

Education have received a planning grant for the creation of a professional development program in mathematics for middle school special education teachers to strengthen skills in the delivery of services to special education students. Initial efforts will analyze the needs of these teachers, including needs that stem from the requirement to prepare students for the High School Exit Exam. The goal of the grant is to define an effective mathematics professional development strategy that meets the needs of special education teachers through traditional face-to-face professional development, online collaboration and online resources that can be used "any time, any place" with the assistance of an online facilitator.

The findings generated by these projects illustrate how the California K–12 High-Speed Network bandwidth can be used to enhance teaching and learning for students' benefit.

John Ittelson, professor and director of The IDEA Lab, and his crew from CSU–Monterey Bay, Jeff McCall and Jotham Fischer-Smith, were on hand to record the presentations and promote information sharing. Video clips of the presentations and other information will be posted at http://www.cpec.ca.gov/.

Contributed by Stephanie Couch, director of statewide initiatives, CENIC. ●

The Teaching and Learning Interchange

continued from page 11

resources and online mentor-moderated discussions. Resources and reflection questions promote adapting the new techniques into present teaching.

Examining the practice of accomplished teachers and discussing classroom management, connections to student demographics and resultant differentiated instruction is the most highly requested form of professional development by teachers. The TLI provides it—via anytime, anywhere access to expert math and science teachers across California. Research has shown that the most effective way to retain teachers is to break the isolation of the classroom and provide support from other professionals. The TLI supports California's efforts to retain math and science teachers by offering each one an online mentor with whom they can confer, regardless of location. As the TLI videos are made available on the gigabit network, teachers across the state will be able view National Board-certified and master teachers' practices while conversing with an accomplished teacher in their field—what a gift to California's teachers and students.

For more information on the TLI, visit www.teachinginterchange.org.

Contributed by Pamela Redmond. ●

Meeting the Challenge of High-Energy Physics

continued from page 16

UltraLight will mark the entry into a new era of global real-time information systems, where all three sets of resources —computational, storage and network—are monitored and tracked to provide efficient, policy-based resource usage. By consolidating with other emerging data-intensive grid systems, UltraLight will drive the next generation of grid and monitoring system developments, as well as new modes of collaborative work. UltraLight paves the way for more flexible, efficient data sharing by scientists in many countries and could be a key factor in the next round of physics discoveries awaiting us at the high-energy frontier. Closer to home, it offers profound implications for integrating into our daily lives information sharing and on-demand audiovisual collaboration at a quality previously unimaginable.

Footnotes

1. The LHCNet is the transatlantic 10 Gbps backbone connecting CERN in Geneva to Chicago, Illinois. http://www.datatag.org

2. NLR is an initiative of U.S. research universities and private technology companies to provide a national infrastructure for research and experimentation in networking technologies and applications. http://www.nlr.net/

3. UCAID is a consortium of 206 universities working in partnership with industry and government to develop and deploy advanced network applications and technologies. http://www.internet2.edu

4. StarLight is a high-performance network exchange for many worldwide research and educational wide-area networks.

5. TransLight is a global, experimental networking initiative that supports prototypes of aggressive e-science applications. http://www.startap.net/translight/

6. NetherLight is an advanced optical infrastructure with international connectivity. http://www.surfnet.nl/innovatie/netherlight/

7. UKlight is a national facility to support projects working on developments towards optical networks. http://www.uklight.ac.uk/

8. The mission of AMPATH is to serve as the pathway for research and education networking in the Americas and to be the international exchange point for Latin America and the Caribbean research and education networks. http://www.ampath.fiu.edu/

9. CA*Net4 embodies a true "customer-empowered network" by placing dynamic allocation of network resources in the hands of end users. http://www.canarie.ca/canet4/

10. FAST is a congestion-control algorithm that improves TCP performance in high-speed networks. It's based on a broader theoretical effort toward robust and stable ultra-scale networking. http://netlab.caltech.edu/FAST/index.html

11. MonALISA is an advanced distributed monitoring system. http://www.monalisa.org

12. Grid Analysis Environment. http://pcbunn.cithep.caltech.edu/GAE/GAE.htm

For more information on the UltraLight project, see http://ultralight.caltech.edu.

Contributed by Sylvain Ravot, Harvey Newman, Julian Bunn, Iosif Legrand and Frank an Lingen. ●

Corporation for Education Network Initiatives in California
5757 Plaza Drive, Suite 205, Cypress, CA 90630
714.220.3400 • www.cenic.org • info@cenic.org

CENIC Board Associates

California Community Colleges

Alameda
Allan Hancock
American River
Antelope Valley
Bakersfield
Barstow
Butte
Cabrillo
Cañada
Canyons
Cerritos
Cerro Coso
Chabot
Chaffey
Citrus
Coastline
Columbia
Compton
Contra Costa
Copper Mountain
Cosumnes River
Crafton Hills
Cuesta
Cuyamaca
Cypress
De Anza
Desert
Diablo Valley
East Los Angeles
El Camino
Evergreen Valley
Feather River
Folsom Lake
Foothill
Fresno City
Fullerton
Gavilan
Glendale
Golden West
Grossmont
Hartnell
Imperial Valley
Irvine Valley
Lake Tahoe
Laney
Las Positas
Lassen
Lemoore
Long Beach City
Los Angeles City
Los Angeles Harbor
Los Angeles Mission
Los Angeles Pierce
Los Angeles Southwest
Los Angeles Trade Tech
Los Angeles Valley
Los Medanos
Marin
Mendocino
Merced
Merritt
MiraCosta
Mission
Modesto Junior
Monterey Peninsula
Moorpark
Mt. San Antonio
Mt. San Jacinto
Napa Valley
Ohlone
Orange Coast
Oxnard
Palomar
Palo Verde
Pasadena City
Porterville
Redwoods
Reedley
Rio Hondo
Riverside
Sacramento City
Saddleback
San Bernardino Valley
San Diego City
San Diego Mesa
San Diego Miramar
San Francisco City
San Joaquin Delta
San Jose City
San Mateo
Santa Ana
Santa Barbara City
Santa Monica
Santa Rosa
Santiago Canyon
Sequoias
Shasta
Sierra
Siskiyous
Skyline
Solano
Southwestern
Taft
Ventura
Victor Valley
Vista
West Hills
West Los Angeles
West Valley
Yuba

California Institute of Technology

Jet Propulsion Laboratory

California State University System

California Maritime Academy
California Polytechnic State University, San Luis Obispo
California State Polytechnic University, Pomona
CSU Bakersfield
CSU Channel Islands
CSU Chico
CSU Dominguez Hills
CSU Fresno
CSU Fullerton
CSU Hayward
CSU Long Beach
CSU Los Angeles
CSU Monterey Bay
CSU Northridge
CSU Sacramento
CSU San Bernardino
CSU San Marcos
CSU Stanislaus
Humboldt State University
San Diego State University
San Francisco State University
San Jose State University
Sonoma State University

Stanford University

Stanford Linear Accelerator Center
Stanford Medical Center

University of California System

UC Berkeley
UC Davis
UC Irvine
UC Los Angeles
UC Merced
UC Riverside
UC San Diego
UC San Francisco
UC Santa Barbara
UC Santa Cruz

University of Southern California

Health Sciences Campus
Information Sciences Institute
University Park Campus

CENIC Partners

California K–12 System

Encompassing 58 counties and 985 districts

CalREN Associates

University and Community College System of Nevada

CENIC 2005

Pathways to DISCOVERY

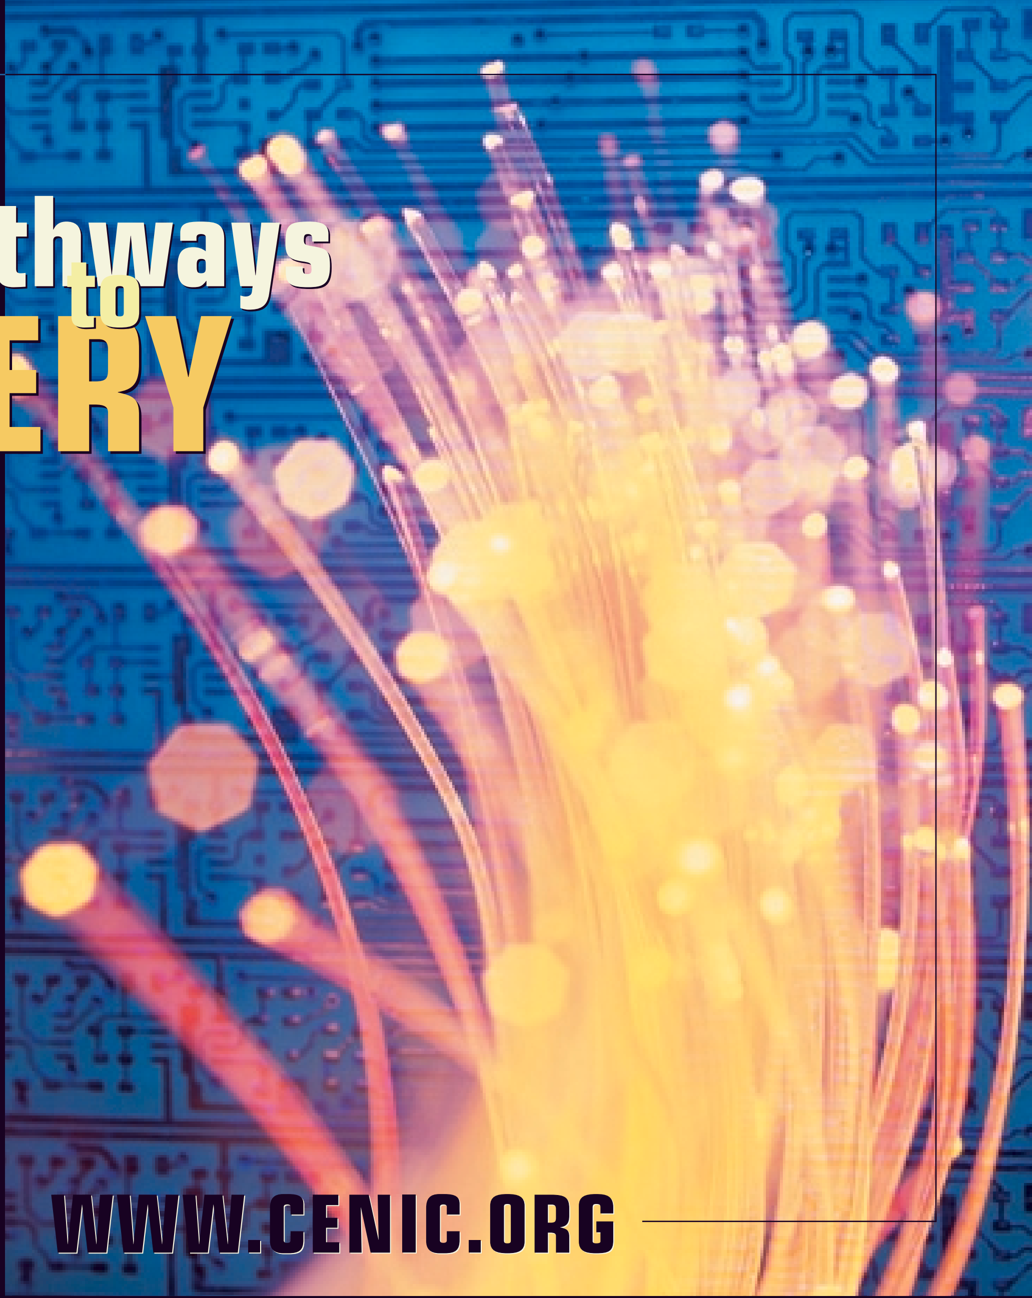

MARCH 7–9, 2005
RITZ-CARLTON MARINA DEL REY
MARINA DEL REY, CALIFORNIA

Corporation for Education Network Initiatives in California (CENIC) invites you to its ninth annual conference, **Pathways to Discovery**. Quite possibly the most important networking and educational event of the year, CENIC 2005 is your opportunity to connect with California's key high-performance networking professionals—educators, researchers, business people and government representatives—and help advance the vision of tomorrow's Internet. This year's conference offers three days of stimulating tracks, influential speakers and a rich array of collaborative panels amidst the engaging backdrop of Marina del Rey, California.

Reserve your hotel and conference space today.

WWW.CENIC.ORG

cenic

Corporation for Education Network Initiatives in California

5757 Plaza Drive, Suite 205
Cypress, CA 90630

Prst Std
U.S. Postage
PAID
Permit #2325
San Diego, CA