Moving Towards Terabit/sec Scientific Dataset Transfers

High Energy Physicists Set New Record for Network Data Transfer

With 110+ Gbps Sustained Rates High Energy Physicists Demonstrate How Long Range Networks Can Be Used Efficiently to Support Leading Edge Science on Four Continents

PORTLAND, Oregon – Building on eight years of record-breaking developments, and on the eve of the restart of the Large Hadron Collider, an international team of high energy physicists, computer scientists, and network engineers led by the California Institute of Technology (Caltech) and partners from the University of Michigan, Fermilab, Brookhaven National Laboratory, CERN, San Diego (UCSD), Florida (UF and FIU), Brazil (Rio de Janeiro State University, UERJ, and two of the State Universities of São Paulo, USP and UNESP), Korea (Kyungpook National University, KISTI), Estonia (NICPB) and the National University of Science and Technology in Pakistan (NUST), joined forces to capture the Bandwidth Challenge award for massive data transfers during the SuperComputing 2009 (SC09) conference.

Caltech's exhibit at SC09 by the High Energy Physics (HEP) group and the Center for Advanced Computing Research (CACR) demonstrated applications for globally distributed data analysis for the Large Hadron Collider (LHC) at CERN, along with Caltech's worldwide collaboration system EVO (Enabling Virtual Organizations; http://evo.caltech.edu) developed with UPJS in Slovakia, and its global network and grid monitoring system MonALISA (http://monalisa.caltech.edu) as well as its Fast Data Transfer application (http://monalisa.caltech.edu) developed in collaboration with the Polytechnica University (Bucharest). The CACR team also showed near-real-time simulations of earthquakes in the Southern California region, experiences in time-domain astronomy with Google Sky, and recent results in multi-physics multi-scale modeling.

New Records

The focus of the exhibit was the HEP team's record-breaking demonstration of storage-to-storage data transfer over wide area networks from two racks of servers and a network switch-router on the exhibit floor. The high-energy physics team's demonstration "Moving Towards Terabit/sec Transfers of Scientific Datasets: The LHC Challenge" achieved a bi-directional peak throughput of 119 gigabits per second (Gbps) and a data flow of more than 110 Gbps that could be sustained indefinitely among clusters of servers on the show floor and at Caltech, Michigan, San Diego, Florida, Fermilab, Brookhaven, CERN, Brazil, Korea, and Estonia.

Following the Bandwidth Challenge the team continued its tests and demonstrated a world-record data transfer between the Northern and Southern hemispheres, sustaining 8.26 Gbps on each of two 10 Gbps links linking São Paulo and Miami.

By setting new records for sustained data transfer among storage systems over continental and transoceanic distances, using simulated LHC datasets, the HEP team demonstrated its readiness to enter a new era in the use of state of the art cyber-infrastructure to enable physics discoveries at the high energy frontier, while demonstrating some of the groundbreaking tools and systems they have developed to enable a global collaboration of thousands of scientists located at 350 universities and laboratories in more than 100 countries to make the next round of physics discoveries.

Advanced Networks, Servers and State of the Art Applications

The record-setting demonstrations were made possible through the use of fifteen 10 Gbps links to SC09 provided by SCinet together with National Lambda Rail (11 links including 6 dedicated links to Caltech) and CENIC, Internet2 (2 links), ESnet, and Cisco. The Caltech HEP team used its dark fiber connection to Los Angeles provided by Level3 and a pair of DWDM optical multiplexers provided by Ciena Corporation to light the fiber with a series of 10G wavelengths to and from the Caltech campus in Pasadena. Ciena also supported a portion of the Caltech traffic with a single serial 100G wavelength running into the SC09 conference from the Portland Level3 PoP, operating alongside other links into SC09 from Portland. Onward connections to the partner sites included links via Esnet and Internet2 to UCSD, FLR to U. Florida as well as FIU and Brazil, MiLR to Michigan, Starlight and USLHCNet to CERN, Ampath together with RNP and ANSP to Brazil, GLORIAD and KREONet to Korea, and Internet2 and GEANT3 to Estonia.

The network equipment at the Caltech booth was a single heavily populated Nexus 7000 series switch-router provided by Cisco, and a large number of 10 gigabit Ethernet server interface cards provided by Myricom. The server equipment on the show floor included five widely available Supermicro 32 core servers using Xeon quad core processors with 12 Seagate SATA disks each, and 18 Sun Fire X4540 servers each with 12 cores and 48 disks provided by Sun Microsystems.

One of the features of next generation networks supporting the largest science programs, notably the LHC experiments, is the use of dynamic circuits with bandwidth guarantees crossing multiple network domains. The Caltech team at SC09 used Internet2's recently announced ION service, developed together with ESnet, GEANT and in collaboration with US LHCNet, to create a dynamic circuit between Portland and CERN as part of the bandwidth challenge demonstrations.

One of the key elements in this demonstration was Fast Data Transfer (FDT), an open source Java application developed by Caltech in close collaboration with Polytehnica University in Bucharest. FDT runs on all major platforms and uses the NIO libraries to achieve stable disk reads and writes coordinated with smooth data flow using TCP across long-range networks. The FDT application streams a large set of files across an open TCP socket, so that a large data set composed of thousands of files, as is typical in high energy physics applications, can be sent or received at full speed, without the network transfer restarting between files. FDT can work on its own, or together with Caltech's MonALISA system, to dynamically monitor the capability of the storage systems as well as the network path in real-time, and send data out to the network at a moderated rate that achieves smooth data flow across long range networks.

Since it was first deployed at SC06, FDT has been shown to reach sustained throughputs among storage systems at 100% of network capacity where needed, in production use, including among systems on different continents. FDT also achieved a smooth bidirectional throughput of 191 Gbps (199.90 Gbps peak) using an optical system carrying an OTU-4 wavelength over 80 km provided by CIENA last year at SC08.

Another new aspect of the HEP demonstration was large scale data transfers among multiple file systems widely used in production in the LHC community, with several hundred Terabytes per site. This included two recently installed instances of the open source file system Hadoop, where in excess of 9.9 Gbps was read from Caltech on one 10 Gbps link, and up to 15 Gbps was read on shared ESnet and NLR links -- a level just compatible with the production traffic on the same links. The high throughput was achieved through the use of a new FDT/Hadoop adaptor-layer written by NUST in collaboration with Caltech.

Lessons Learned: Towards A Compact Terabit/sec Facility

The SC09 demonstration also achieved its goal of clearing the way to Terabit/sec data transfers. The 4-way Supermicro servers at the Caltech booth, each with four 10GE Myricom interfaces, provided 8.3Gbps of stable throughput each, reading or writing on 12 disks, using FDT. A system capable of one Terabit/sec (Tbps) to or from storage could therefore be built today in just six racks at relatively low cost, while also providing 3840 processing cores and 3 Petabytes of disk space, which is comparable to the larger LHC centers in terms of computing and storage capacity.

An important ongoing theme of SC09, including at the Caltech booth where the EVOGreen initiative (www.evogreen.org) was highlighted, was the reduction of our carbon footprint through the use of energy-efficient information technologies. A particular focus is the use of systems with a high ratio of computing and I/O performance to energy consumption, for which the SC09 Storage Challenge entry "*LowPower Amdahl-Balanced Blades for Data Intensive Computing"* by Szalay *et al.* is a notable example. In the coming year, in preparation for SC10 in New Orleans, the HEP team will be looking into the design and construction of compact systems with lower power and cost that are capable of delivering data at several hundred Gbps, aiming to reach 1 Tbps by 2011 when multiple 100 Gbps links into SC11 may be available.

The LHC Program: CMS and ATLAS

The two largest physics collaborations at the LHC, CMS and ATLAS, each encompassing more than 2,000 physicists, engineers and technologists from 180 universities and laboratories, are about to embark on a new round of exploration at the frontier of high energies. When the LHC experiments begin to take collision data in a new energy range over the next few months, new ground will be broken in our understanding of the nature of matter and space-time and in the search for new particles. In order to fully exploit the potential for scientific discoveries during the next year, more than 100 petabytes (10^{17} bytes) of data will be processed, distributed, and analyzed using a global grid of 300 computing and storage facilities located at laboratories and

universities around the world, rising to the exabyte range (10^{18} bytes) during the following years.

The key to discovery is the analysis phase, where individual physicists and small groups located at sites around the world repeatedly access, and sometimes extract and transport multi-terabyte data sets on demand from petabyte data stores, in order to optimally select the rare "signals" of new physics from potentially overwhelming "backgrounds" from already-understood particle interactions. The HEP team hopes that the demonstrations at SC09 will pave the way towards more effective distribution and use for discoveries of the masses of LHC data.

Quotes on the significance of the demonstrations:

"By sharing our methods and tools with scientists in many fields, we hope that the research community will be well-positioned to further enable their discoveries, taking full advantage of current networks, as well as next-generation networks with much greater capacity as soon as they become available. In particular, we hope that these developments will afford physicists and young students throughout the world the opportunity to participate directly in the LHC program, and potentially to make important discoveries."

-- Harvey Newman, Caltech professor of physics, head of the HEP team and co-lead of US LHCNet, and chair of the US LHC Users Organization

"The efficient use of high-speed networks to transfer large data sets is an essential component of CERN's LHC Computing Grid (LCG) infrastructure that will enable the LHC experiments to carry out their scientific missions."

-- David Foster, Deputy IT Department Head, co-lead of US LHCNet and former Head of Communications and Networking at CERN

"We continue to demonstrate the state of the art in realistic, worldwide deployment of distributed, data-intensive applications capable of effectively using and coordinating high-performance networks… Our distributed agent-based autonomous system is used to dynamically discover network and storage resources, and to monitor, control, and orchestrate efficient data transfers among hundreds of computers, as well as tens of millions of jobs per year, and the complete topology of dynamic circuits in networks such as US LHCNet."

-- Iosif Legrand, senior software and distributed system engineer at Caltech, the technical coordinator of the MonALISA and FDT projects. "

"This achievement is an impressive example of what a focused network and storage system effort can accomplish. It is an important step towards the goal of delivering a

highly capable end-to-end network-aware system and architecture that meet the needs of next-generation e-science."

-- *Shawn McKee, research scientist in the University of Michigan department of physics and leader of the UltraLight network technical group*

"The impressive capability of dynamically setting up the many light paths used in this demonstration in such a short time frame, spanning three continents and providing guaranteed bandwidth channels for applications requiring them, together with the efficient use of the provisioned bandwidth by the data transfer applications, shows the high potential in circuit network services. The light path setup among USLHCNet, Surfnet, CANARIE, TransLight/StarLight, ESnet SDN, and Internet2 ION, and using the MANLAN, Starlight, and Netherlight exchange points, took only days to accomplish (minutes in the case of SDN and DCN dynamic circuits). It shows how the network can already today be used as a dedicated resource in data intensive research and other fields, and demonstrates how applications can make best use of this resource basically on demand."

-- *Artur Barczyk, lead engineer of US LHCNet and head of the SC09 and SC08 network engineering teams*

“Participation in this year’s bandwidth challenge has had a tremendous impact at Florida International University, not only in commissioning our new 10 GigE campus network infrastructure but also by creating a unique opportunity for our students to participate in a challenging, rewarding and above all exciting experience.”

-- *Jorge Luis Rodriguez, Assistant Professor of Physics at FIU*

"When you combine this network-storage technology, including its cost profile, with the remarkable tools that Harvey Newman's networking team has produced, I think we are well positioned to address the incredible infrastructure demands that the LHC experiments are going to make on our community worldwide."

-- *Paul Sheldon, Professor of Physics at Vanderbilt and leader of REDDNet*

“Caltech and partners have not only set new records for massive data transfer across wide-area networks, but also proven that geographic distance need not be a barrier for even the most demanding, high-bandwidth applications. National LambdaRail is very proud to have provided 11 10-GE circuits, 6 of which were dedicated, for the collaborators of this breakthrough demonstration and we congratulate Caltech and partners on their SC09 Bandwidth Challenge award."

-- Glenn *Ricart, President and CEO of National LambdaRail*

Acknowledgements

The demonstration and the developments leading up to it were made possible through the strong support of the partner network organizations mentioned, the U.S. Department of Energy Office of Science, and the National Science Foundation, in cooperation with the

funding agencies of the international partners, through the following grants: US LHCNet (DOE DE-FG02-05-ER41359), UltraLight (NSF PHY-0427110), DISUN (NSF PHY-0533280), CHEPREO/WHREN-LILA (NSF PHY-0312038/OCI-0441095 and FAPESP Projeto 04/14414-2), the NSF-funded PLaNetS (NSF PHY-0622423) project, an NSF travel specifically for SC09, and the US LHC Program funded jointly by DOE and NSF.

Further information about the demonstration may be found at:
http://supercomputing.caltech.edu

About Caltech: With an outstanding faculty, including five Nobel laureates, and such off-campus facilities as the Jet Propulsion Laboratory, Palomar Observatory, and the W. M. Keck Observatory, the California Institute of Technology is one of the world's major research centers. The Institute also conducts instruction in science and engineering for a student body of approximately 900 undergraduates and 1,300 graduate students who maintain a high level of scholarship and intellectual achievement. Caltech's 124-acre campus is situated in Pasadena, California, a city of 135,000 at the foot of the San Gabriel Mountains, approximately 30 miles inland from the Pacific Ocean and 10 miles northeast of the Los Angeles Civic Center. Caltech is an independent, privately supported university, and is not affiliated with either the University of California system or the California State Polytechnic universities. http://www.caltech.edu.

About CACR: The mission of the Center for Advanced Computing Research (CACR) is to ensure that Caltech is at the forefront of computational science and engineering. CACR provides an environment that cultivates multidisciplinary collaborations and its researchers take an applications-driven approach and currently work with Caltech research groups in aeronautics, applied mathematics, astronomy, biology, engineering, geophysics, materials science, and physics. Center staff have expertise in data-intensive scientific discovery, physics-based simulation, scientific software engineering, visualization techniques, novel computer architectures, and the design and operation of large-scale computing facilities. http://www.cacr.caltech.edu/.

About CERN: CERN, the European Organization for Nuclear Research, has its headquarters in Geneva. At present, its member states are Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, the Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland, and the United Kingdom. Israel, Japan, the Russian Federation, the United States of America, Turkey, the European Commission, and UNESCO have observer status. For more information, see http://www.cern.ch.

About the University of Michigan: The University of Michigan, with its size, complexity, and academic strength, the breadth of its scholarly resources, and the quality of its faculty and students, is one of America's great public universities and one of the world's premier research institutions. The university was founded in 1817 and has a total enrollment of 54,300 on all campuses. The main campus is in Ann Arbor, Michigan, and has 39,533 students (fall 2004). With over 600 degree programs and $739M in FY05

research funding, the university is one of the leaders in innovation and research. For more information, see http://www.umich.edu.

About the University of Florida: The University of Florida (UF) is a major public land-grant research university; the state's oldest, largest and most comprehensive, and among the most academically diverse in the nation. UF has a long history of established programs in international education, research and service with $574 million in sponsored research grants in 2008. With an enrollment approaching 50,000 students annually, UF is home to 16 colleges and more than 150 research centers and institutes. The Institute for High Energy Physics and Astrophysics (IHEPA) of the UF Physics Department and the UF High Performance Computing Center (UF HPC) have constructed and now operate one of only seven CMS Tier 2 centers in US. IHEPA is playing a leading role in the development of Grid technologies for distributed computing and data analysis for the CMS experiment.

About Florida International University: Florida International University is one of Florida's State University System institutions and the premier public research university in the South Florida metropolitan area, the country's fourth largest. The university is located in Miami-Dade County with campuses in West and North Miami. The university with its large and diverse student body is the largest minority serving institution in the nation and is ranked first in awarding STEM degrees to underrepresented minorities. The university was founded in 1972 with only 5,667 students. In the four short decades since the university has grown phenomenally and now boasts a total enrollment of more than 39,000 students; 60 percent of which are of Hispanic descent. With an increasing emphasis on graduate research, the university has the Carnegie Foundation's highest ranking and in 2008, received over $100 million in external contracts and grants. http://www.fiu.edu

About UCSD: Nestled along the Pacific Ocean on 1,200 acres of coastal woodland, UC San Diego is a powerful magnet for those seeking a fresh approach to education and research. UC San Diego is dedicated to the advancement of knowledge through excellence in education and research at the undergraduate, graduate, professional school and postdoctoral levels. The campus is committed to community engagement, public service and industry partnerships in order to advance the health and well-being of our region, state, nation and the world. Our academic community of world-renowned faculty, bright students and dedicated staff is characterized by a culture of interdisciplinary collaboration and innovation which spans the globe. UC San Diego's high academic rankings reflect its status as one of the top institutions in the nation for higher education and research.

About Fermilab: Fermi National Accelerator Laboratory (Fermilab) is a national laboratory funded by the Office of Science of the U.S. Department of Energy, operated by Fermi Research Alliance, LLC. Experiments at Fermilab's Tevatron, the world's highest-energy particle accelerator, generate petabytes data per year, and involve large, international collaborations with requirements for high-volume data movement to their home institutions. It is also the western hemisphere Tier-1 data host for the upcoming

CMS experiment at the HC. The laboratory actively works to remain on the leading edge of advanced wide-area network technology in support of its science collaborations.

About Brookhaven National Lab: One of ten national laboratories overseen and primarily funded by the Office of Science of the U.S. Department of Energy (DOE), Brookhaven National Laboratory conducts research in the physical, biomedical, and environmental sciences, as well as in energy technologies and national security. Brookhaven Lab also builds and operates major scientific facilities available to university, industry, and government researchers. Brookhaven is operated and managed for DOE's Office of Science by Brookhaven Science Associates, a limited-liability company founded by the Research Foundation of the State University of New York, for and on behalf of Stony Brook University, the largest academic user of Laboratory facilities; and Battelle Memorial Institute, a nonprofit, applied science and technology organization. Visit Brookhaven Lab's electronic newsroom for links, news archives, graphics, and more (http://www.bnl/gov/newsroom), or follow Brookhaven Lab on Twitter (http://twitter.com/BrookhavenLab).

About the Politehnica University (Bucharest, Romania): Founded in 1818, Politehnica University of Bucharest (UPB: http://www.upb.ro) is the largest and the best Technical University in Romania. UPB is a full member of several international organizations such as CESAER, EUA and AUF, and has bilateral co-operation agreements with similar universities, mainly in Europe, the U.S., Singapore and Japan. UPB also participates in projects funded by NATO and the EU 6th and 7th Frameworks. 26,000 undergraduate, masters and Ph. D students are enrolled at UPB, including more than 1,500 in diverse areas of Computational Science and Engineering (CSE). The Romanian National Center for Information Technology (NCIT) is part of UPB and is run by the CSE Department. UPB has extensive experience in monitoring distributed resources, in projects such as MonALISA, a fully distributed monitoring system based on autonomous, self-describing agent-based subsystems which has been developed over the last seven years by Caltech and UPB. The UPB team also has been involved in the EU-NCIT-EU IST Excellency project, focused on Grid computing and Collaborative work, and the FP7 project SENSEI. UPB also benefits from its results in the FP7 P2P-Next project, such as unified P2P technologies for live streaming and progressive downloading.

About UERJ (Rio de Janeiro): Founded in 1950, the Rio de Janeiro State University (UERJ; http//www.uerj.br) ranks among the ten largest universities in Brazil, with more than 23,000 students. UERJ's five campuses are home to 22 libraries, 412 classrooms, 50 lecture halls and auditoriums, and 205 laboratories. UERJ is responsible for important public welfare and health projects through its centers of medical excellence, the Pedro Ernesto University Hospital (HUPE) and the Piquet Carneiro Day-care Policlinic Centre, and it is committed to the preservation of the environment. The UERJ High Energy Physics group includes 15 faculty, postdoctoral, and visiting Ph.D. physicists and 12 Ph.D. and master's students, working on experiments at Fermilab (D0) and CERN (CMS). The group has constructed a Tier2 center to enable it to take part in the Grid-based data analysis planned for the LHC, and has originated the concept of a Brazilian "HEP Grid," working in cooperation with USP and several other universities in Rio and São Paulo.

About UNESP (São Paulo): Created in 1976 with the administrative union of several isolated institutes of higher education, the São Paulo State University, UNESP, has 39 institutes in 23 different cities in the State of São Paulo. The university has 33,500 undergraduate students in 168 different courses and almost 13,000 graduate students. UNESP has just inaugurated the Center for Scientific Computing with the goal of empowering research groups with high performance computing, storage and networking resources. The new state-of-the-art data center houses the GridUNESP main cluster and the SPRACE cluster.

About SPRACE: The São Paulo Regional Analysis Center (SPRACE) is a Worldwide LHC Computing Grid Tier-2 center operating in association with the Open Science Grid. The SPRACE researchers are members of the Fermilab´s DZero experiment since 2004 and of the CERN's CMS Collaboration. SPRACE has been leveraging competences in different research areas by sharing the technical expertise generated by High Energy Physics experiments, including high speed networks, high performance computing, and grid computing architectures. SPRACE has inspired the GridUNESP project which has deployed the first Campus Grid in Latin America.

About GridUNESP: UNESP has deployed a distributed computational system which is the largest Campus Grid initiatives in Latin America, with computing resources located on seven different campuses. The GridUNESP computational infrastructure includes almost 400 servers, over 200 terabytes of storage space and an advanced networking infrastructure for inter-cluster connection. GridUNESP has established a formal partnership with the Open Science Grid Consortium and is employing the OSG middleware stack to integrate its computational resources and share them with other research and education institutions worldwide.

About NUST/SEECS: The School of Electrical Engineering and Computer Science of NUST, the National University of Sciences and Technology, ranks today among the premier engineering institutions in Pakistan . It possesses state-of-the-art equipment and prides on its faculty, a team of highly capable and dedicated professionals. SEECS is also reputed abroad for its collaborative research linkages with Caltech, USA; CERN, Geneva; Stanford (SLAC), USA; and EPFL, Switzerland, just to name a few. NUST continually seeks to stay in forefront as a centre of academic and research excellence. SEECS also is engaged in close interaction with indigenous industrial entrepreneurs in IT, electronics and communication engineering through SEECS-Corporate Advisory councils. http://www.seecs.edu.pk

About Kyungpook National University (Daegu): Kyungpook National University is one of leading universities in Korea, especially in physics and information science. The university has 13 colleges and 9 graduate schools with 24,000 students. It houses the Center for High Energy Physics (CHEP) in which most Korean high-energy physicists participate. CHEP (chep.knu.ac.kr) was approved as one of the designated Excellent Research Centers supported by the Korean Ministry of Science.

About GLORIAD: GLORIAD (GLObal RIng network for Advanced application development) is the first round-the-world high-performance ring network jointly established by Korea, the United States, Russia, China, Canada, the Netherlands, and the Nordic countries, with optical networking tools that improve networked collaboration with e-Science and Grid applications. It is currently constructing a dedicated lightwave link connecting the scientific organizations in GLORIAD partners. See http://www.gloriad.org/.

About KISTI: KISTI is a specialized institute providing STI (Science and Technology Information) services based on national supercomputing center and advanced research networks (KREONET, GLORIAD-KR and KRLight) to promote global competitiveness in science and technology by actively challenging the rapidly changing world paradigm. KRLight supports end-to-end dedicated lightpath provisioning for high end applications such as HEP as a GLIF Open Lightpath Exchange of Korea. KISTI aims to be the world leader in science and technology.

About National LambdaRail: Owned by the U.S. research and education community and dedicated to serving the needs of researchers and research groups, NLR is the innovation platform for a wide range of academic disciplines and public-private partnerships. NLR's coast-to-coast, high-performance network infrastructure offers unrestricted usage and bandwidth, a choice of cutting-edge network services and applications, and customized service for individual researchers and projects. For more information, please visit www.nlr.net.

About Internet2: Internet2 is the foremost U.S. advanced networking consortium. Led by the research and education community since 1996, Internet2 provides both leading-edge network capabilities and unique partnership opportunities that together facilitate the development, deployment and use of revolutionary Internet technologies. Internet2 brings the U.S. research and academic community together with technology leaders from industry, government and the international community to undertake collaborative efforts that have a fundamental impact on tomorrow's Internet. For more information: http://www.internet2.edu

About StarLight: StarLight is an advanced optical infrastructure and proving ground for network services optimized for high-performance applications. Operational since summer 2001, StarLight is a 1 GE and 10 GE switch/router facility for high-performance access to participating networks and also offers true optical switching for wavelengths. StarLight is being developed by the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago (UIC), the International Center for Advanced Internet Research (iCAIR) at Northwestern University, and the Mathematics and Computer Science Division at Argonne National Laboratory, in partnership with Canada's CANARIE and the Netherlands' SURFnet. STAR TAP and StarLight are made possible by major funding from the U.S. National Science Foundation to UIC. StarLight is a service mark of the Board of Trustees of the University of Illinois. See www.startap.net/starlight.

About the Florida LambdaRail: Florida LambdaRail LLC (FLR) is a Florida limited liability company formed by member higher education institutions to advance optical research and education networking within Florida. Florida LambdaRail is a high-bandwidth optical network that links Florida's research institutions and provides a next-generation network in support of large-scale research, education outreach, public/private partnerships, and information technology infrastructure essential to Florida's economic development. For more information: http://www.flrnet.org.

About CENIC: California's education and research communities leverage their networking resources under CENIC, the Corporation for Education Network Initiatives in California, to fulfill the needs of their faculty, staff, and students for cost-effective, high-bandwidth networking. CENIC designs, implements, and operates the 2,900 mile California Research & Education Network (CalREN) backbone, a high-capacity three-tiered DWDM network composed of CENIC-owned optical fiber and several hundred optical components, to which the vast majority of the state's K-20 educational institutions in all 58 California counties are connected via leased circuits from telecom carriers or fiber-optic cable. CENIC is governed by its member institutions. Representatives from these institutions donate their expertise through participation in committees designed to ensure that CENIC is managed effectively and efficiently, and to support the continued evolution of the network as technology advances. For more information visit www.cernic.org

About ESnet: The Energy Sciences Network (ESnet; www.es.net) is a high-speed network serving thousands of Department of Energy scientists and collaborators worldwide. A pioneer in providing high-bandwidth, reliable connections, ESnet enables researchers at national laboratories, universities, and other institutions to communicate with each other using the collaborative capabilities needed to address some of the world's most important scientific challenges. Managed and operated by the ESnet staff at Lawrence Berkeley National Laboratory, ESnet provides direct high-bandwidth connections to all major DOE sites, multiple cross connections with Internet2/Abilene, and connections to Europe via GEANT and to Japan via SuperSINET, as well as fast interconnections to more than 100 other networks. Funded principally by DOE's Office of Science, ESnet services allow scientists to make effective use of unique DOE research facilities and computing resources, independent of time and geographic location.

About AMPATH: Florida International University's Center for Internet Augmented Research and Assessment (CIARA) has developed an international, high-performance research connection point in Miami, Florida, called AMPATH (AMericasPATH; www.ampath.fiu.edu). AMPATH's goal is to enable wide-bandwidth digital communications between U.S. and international research and education networks, as well as a variety of U.S. research programs in the region. AMPATH in Miami acts as a major international exchange point (IXP) for the research and education networks in South America, Central America, Mexico, and the Caribbean. The AMPATH IXP is home for the WHREN-LILA high-performance network link connecting Latin America to the U.S., funded by the NSF, the Academic Network of Sao Paulo, and the Rede Nacional de Ensino e Pesquisa (RNP).

About the Academic Network of São Paulo (ANSP): ANSP unites São Paulo's University networks with Scientific and Technological Research Centers in São Paulo, and is managed by the State of São Paulo Research Foundation (FAPESP). The ANSP Network is another example of international collaboration and exploration. Through its connection to WHREN-LILA, all of the institutions connected to ANSP will be involved in research with U.S. universities and research centers, offering significant contributions and the potential to develop new applications and services. This connectivity with WHREN-LILA and ANSP will allow researchers to enhance the quality of current data, inevitably increasing the quality of new scientific developments. http://www.ansp.br.

About Rede Nacional de Ensino e Pesquisa (**RNP):** RNP, the National Education and Research Network of Brazil, is a not-for-profit company that promotes the innovative use of advanced networking, with the joint support of the Ministry of Science and Technology and the Ministry of Education. In the early 1990s, RNP was responsible for the introduction and adoption of Internet technology in Brazil. Today, RNP operates a nationally deployed multi-gigabit Ipê network used for collaboration and communication in research and education throughout the country, reaching all 26 states and the Federal District, and provides both commodity and advanced research Internet connectivity to more than 600 universities, research centers, and technical colleges. Additionally, together with the telecommunications industry research center, CPqD, RNP operates the experimental GIGA network, which provides high bandwidth connectivity to selected laboratories in the states of Rio de Janeiro and São Paulo, including the HEPGrid facility at UERJ. http://www.rnp.br

About Southern Light: Southern Light (SOL), the only GLIF Open Lightpath Exchange (GOLE) in the Southern Hemisphere, is located in the city of São Paulo, and is operated jointly by ANSP and RNP. SOL provides international connectivity through two 10G links to the AMPATH exchange point in Miami for institutions connected to RNP's Ipê and GIGA networks, and to FAPESP's KyaTera and ANSP networks. See http://wiki.glif.is/index.php/SouthernLight

About Intel: Intel, the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at www.intel.com/pressroom and blogs.intel.com.

About Sun Microsystems: Sun Microsystems develops the technologies that power the global marketplace. Guided by a singular vision -- "The Network is the Computer" -- Sun drives network participation through shared innovation, community development and open source leadership. Sun can be found in more than 100 countries and on the Web at http://sun.com.

About Myricom: Founded in 1994, Myricom created Myrinet, the pioneering High Performance Computing (HPC) interconnect technology used in many thousands of computing clusters in more than 50 countries. With its fourth generation of networking products, Myri-10G, Myricom achieved a convergence at 10-Gigabit data rates between its low-latency Myrinet technology and mainstream Ethernet. These Myri-10G products

continue to be used with Myrinet protocols for HPC clusters, but the majority of Myricom's sales today are for 10-Gigabit Ethernet applications with demanding latency or throughput requirements. Systems with Myricom networking are supplied by OEM computer companies and leading cluster integrators worldwide. See www.myri.com.

About the National Science Foundation: The NSF is an independent federal agency created by Congress in 1950 "to promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense...." With an annual budget of about $5.5 billion, it is the funding source for approximately 20 percent of all federally supported basic research conducted by America's colleges and universities. In many fields such as mathematics, computer science, and the social sciences, NSF is the major source of federal backing.

About the DOE Office of Science: DOE's Office of Science is the single largest supporter of basic research in the physical sciences in the nation and ensures U.S. world leadership across a broad range of scientific disciplines. The Office of Science also manages 10 world-class national laboratories with unmatched capabilities for solving complex interdisciplinary problems, and it builds and operates some of the nation's most advanced R&D user facilities, located at national laboratories and universities. These facilities are used by more than 19,000 researchers from universities, other government agencies, and private industry each year.

Acknowledgements

The demonstration and the developments leading up to it were made possible through the strong support of the partner network organizations mentioned, the U.S. Department of Energy Office of Science and the National Science Foundation, in cooperation with the funding agencies of the international partners, through the following grants: US LHCNet (DOE DE-FG02-05-ER41359), WAN In Lab (NSF EIA-0303620), UltraLight (NSF PHY-0427110), DISUN (NSF PHY-0533280), CHEPREO/WHREN-LILA (NSF PHY-0312038/OCI-0441095 and 0802184), travel grant from NSF specifically for SC09 (NSF PHY-0956884), FAPESP (São Paulo) Projeto 04/14414-2), as well as the NSF-funded PLaNetS and FAST TCP projects, and the US LHC Research Program funded jointly by DOE and NSF.

Contacts:

Harvey B. Newman, Caltech High Energy Physics: newman@hep.caltech.edu
(626) 395-6656

Jon Weiner, Caltech Media Relations: jweiner@caltech.edu (626) 395-3226