Enabling Scientific Discoveries with LHC Data Distribution Over 100 Gigabit Networks

High Energy Physicists Smash Previous Records for Network Data Transfer

Physicists Demonstrate New Methods for Efficient Use of Long Range Networks to Support Leading Edge Science at Breakthrough Data Rates

SALT LAKE CITY, Utah – During the SuperComputing 2012 (SC12) conference November 12-16, an international team of high energy physicists, computer scientists, and network engineers led by the California Institute of Technology (Caltech), the University of Victoria, and the University of Michigan, together with Brookhaven National Lab, Vanderbilt and other partners, smashed their previous records for data transfers using the latest generation of wide area network circuits.

With three 100 gigabit/sec (100 Gbps) wide area network circuits set up by the SCinet, Internet2, CENIC, CANARIE and BCnet, Starlight and US LHCNet network teams, and servers at each of the sites with 40 gigabit Ethernet (40GE) interfaces, the team reached a record transfer rate of 339 Gbps between Caltech, the University of Victoria Computing Center in British Columbia, the University of Michigan, and the Salt Palace Convention Center in Utah. This nearly doubled last year's overall record, and eclipsed the record for a bidirectional transfer on a single link with a data flow of 187 Gbps between Victoria and Salt Lake.

Several other records achieved November 14-15 included: a record overall storage to storage rate using the three links of 187 Gbps, a unidirectional transfer between storage systems in Victoria and Salt Lake on one link of 96 Gbps, an 80 Gbps transfer from Caltech to a single server with two 40GE interfaces at Salt Lake with nearly 100% use of the servers' interfaces at both ends, and a transfer using Remote Data Memory Access (RDMA) over Ethernet between Pasadena and Salt Lake that sustained 75 Gbps with a CPU load on the servers of only 5%. A total of 3.8 Petabytes was transferred over the three days of the conference exhibit, including 2 Petabytes on the last day.

The latest generation of servers based on the PCI Express 3.0 standard and equipped with line-rate 40GE interface cards as well as RAID arrays with solid state disks (SSDs) allowed the team to reach a stable throughput of 38 Gbps from disk to disk over long distances between a pair of two-rack-unit (2U) servers. This tripled the throughput per server achieved with earlier versions of these systems in 2011. Servers equipped with solid state storage cards and two network interfaces reached very close to 80 Gbps in a single server.

While the wide area records were being set, another 50 Gbps flow was established over a 100 Gbps link between the Caltech and Vanderbilt booths at SC12, supported by Padtec, previewing a demonstration between Caltech, Florida and Brazil in the coming months. The peak flow to and from the Caltech booth thus reached 389 Gbps.

Impact on Data Intensive Science: Terabit/sec Rates in Reach

The SC12 demonstrations achieved their goals of clearing the way to the next level of data intensive science. With the methods and configurations shown by the team at SC12, a single rack can now be deployed that is capable of aggregate sustained transfers of more than 1 Terabit/sec, allowing 10 Petabyte datasets to be moved among remote storage systems in a single day. This represents a major milestone expanding the reach of data intensive science, not only in high energy physics, but in many fields that deal with "Big Data": from astrophysics to genomics, meteorology and global climate tracking, connectomics, complex physics simulations, and biological and environmental research. The SC12 demonstrations also show the way towards pushing back the limitations encountered in Internet search, finance and business informatics.

A near term step foreseeable in the coming months is individual servers with 100GE interfaces fully matching a 100 Gbps wide area network link.

The LHC Program: CMS and ATLAS

The two largest physics collaborations at the LHC, CMS and ATLAS, each encompassing more than 3,000 physicists, students, engineers and technologists from 180 universities and laboratories, have embarked on a new round of exploration at the frontier of high energies. The pivotal discovery in July of a new particle at 125 GeV, that may be the long sought Higgs boson linked to mass in the universe, was made possible by the combined efforts of physicists and students using computational and storage facilities at more than 300 sites around the world, who relied on data transfers that reached more than 100 Petabytes during the past year alone.

The LHC program's reach for new discoveries continues to expand. New ground will be broken in the search for new particles, for a fundamental theory to explain the nature of dark matter and physics as it existed in the first moments of the universe, and in our understanding of the nature of matter and space-time. The experiments have already tripled their datasets this year since the July discovery, and have begun to present new results, in a cycle that will be completed next Spring. The LHC and its experiments will then shut down to install upgrades in preparation for renewed running in late 2014 or 2015 at close to the design energy of 14 TeV (compared to 8 TeV at present) that will open new vistas for exploration.

As the LHC experiments continue to take data at increasing collision rates and energy, the stored data volume is expected to rise to the Exabyte range (10^{18} bytes). Transfer rates will continue their exponential growth, spurred on by the ongoing progress in server and network systems, moving from one technology generation to the next, and new data transport methods, as shown by this year's demonstrations by the HEP team at SC12.

State of the Art Applications for Data Intensive Science

Caltech's exhibit at SC12 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 global network and grid monitoring system MonALISA (http://monalisa.caltech.edu). The worldwide collaboration system SeeVogh

(http://www.seeogh.com) developed by Caltech's Enabling Virtual Organizations (EVO) team (http://evo.caltech.edu) was another highlight of the exhibit.

A key element of the exhibit, and the record demonstrations, is Fast Data Transfer (http://monalisa.caltech.edu/FDT), an open source TCP-based Java application developed by the Caltech team in collaboration with Polytehnica University in Bucharest. FDT runs on all major platforms and uses the Java 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 during SC08.

FDT is now being used in production for high throughput transfers in high energy physics as well as in many other data intensive areas, including the Global Environment for Network Innovations (GENI).

Another key element was the FTP 100 software for RDMA-enabled data transfer, which bypasses the operating system kernel to achieve very high throughput over lossless wide area networks with very low CPU loads on the end-systems. This can be very beneficial for a wide range of data-intensive applications, such as those in cloud computing, grid computing, and data centers. The FTP 100 project, part of the DOE ARRA funded 100Gbps network program, aims to provide high-performance data transfer capability to many application domains, including high energy and nuclear physics, climate simulation, bioinformatics, and multi-petabyte data migration among data centers. The successful demonstrations at SC12 will help the FTP 100 team disseminate this new networking and data transfer technology to more big data applications that require throughputs of up to 100 Gbps.

Advanced Networks

The University of Victoria (UVic) with network partners CANARIE, BCNet and Internet 2 established a 100G link between University of Victoria Data Centre and the Salt Palace Convention centre in Salt Lake City Utah. UVic focused on producing maximum disk read and write performance using the smallest amount of hardware possible. UVic was able to produce disk reads above 95 Gbps using only 4 IBM x3650 M4 servers each populated with 16 OCZ Vertex 4 SSDs. An individual IBM server was to read from disk at 38 Gbps and write stably at 24 Gbps. Each machine used a single Mellanox Connect-X3 Ethernet card connected to a Juniper MX 480 with QSFP+ optics.

The connection was carried out of the MX 480 via 100G LR4 CFP Optic to Ciena OME 6500 for Optical Transport over CANARIE and BCNet to Seattle. From Seattle the circuit was transported over the Internet2 100G SDN Infrastructure to Salt Lake City. The established circuit was flawless with zero packet drops allowing for very high sustained throughput.

The Caltech setup at SC12 was built in partnership with several network and storage vendors, including Alcatel Lucent, Juniper, Data Direct Networks, Dell, Fusion-io, Mellanox, OCZ, Padtec and SuperMicro. The wide area network with three 100G links was provided with the support of Internet2, SCinet, Ciena, Cisco, CENIC, CANARIE and BCNet. The 100GE connections were terminated using LR10 optics on Alcatel and Juniper routers, while 40GE server ports were terminated on a Dell Z9000 switch using Dell QSFP SR optics and MTP cables.

The storage appliance provided by Data Direct Networks (DDN) included two SFA-12k controllers connected to five disk enclosures with a total of 288 SAS 15k RPM disks each of 450GB capacity. Eight Lustre Object Storage Servers (OSS) were connected to the DDN controllers through an Infiniband (IB) fabric using Mellanox 56Gbps FDR interconnects. Six Lustre clients were used to read and write over the wide area network; using a second Mellanox network interface on each OSS configured for 40GE.

Fusion-io provided ten ioDrive Duo PCIe storage cards, each having a capacity of 1.2TB, which were installed in two 3U SuperMicro Sandy Bridge Servers. These servers were connected directly to the 40GE ports on the Alcatel Router to avoid any inter-switch congestion. One server with five Fusion IO drives was able to sustain a steady disk write of 80Gbps.

Other storage servers which were part of the overall disk transfers included two 24 bay 2U SuperMicro chasses each housing OCZ Vertex-4 drives chassis and three LSI controllers. Each of these servers had a measured sustained write rate of 32.5 Gbps or just above 4GB/s. A third SuperMicro server with 24 OCZ Vertex-3 drives, and a fourth server with four OCZ PCIe cards, were measured to sustain 18 and 30 Gbps respectively. All of these servers are based on Dual Intel E5 SandyBridge processors, 64GB of RAM and a Mellanox CX-3 VPI network interface card.

Caltech and Vanderbilt partnered with Padtec, an optical switch manufacturer from Brazil to demonstrate a single wave 100GE connection between the two booths on the exhibit floor. At the Vanderbilt booth the client ports from the optical switch were terminated on a Dell Z9000 switch, where another 40GE server was connected for transfers using FDT.

As one of the 100GE End-Site, Caltech in partnership with Cisco, Alcatel Lucent, CENIC and Internet2 established a 100GE link between the CENIC PoP in Los Angeles and the Caltech campus in Pasadena using a Cisco ONS 15454 M6 DWDM Optical platform over a Level3 dark fiber pair. At the CENIC PoP, the 100GE link from the Cisco M6 was connected to Internet2's Advanced Layer2 Services Backbone (AL2S), while at Caltech the 100GE link coming from the M6 was terminated on an Alcatel Lucent SR-12 router hosted in Caltech's CMS Tier2 Center. Three SuperMicro Twin2U servers connected to the Alcatel router using Mellanox 40GE NICs at Caltech were used to read from memory in Pasadena and write on the disk servers in Salt Lake City. Two such servers from

Caltech were able to read at a total speed of 80Gbps (10 GB/s), with sustained disk writes of 10GB/s on a single server using five Fusion-io Duo cards on the SC12 show floor.

The University of Michigan, partnering with ADVA and Juniper, created a 100G path between the ATLAS Great Lakes Tier-2 (AGLT2) Center on the Ann Arbor UM camps and the Internet2 Brocade MLX in Chicago. Internet2 provided transport between Chicago and the SC12 showroom floor in Salt Lake City, and AGLT2 provided the local networking and servers participating in SC12. The 100G wave terminated on a Juniper MX480 situated in the AGLT2 machine room. The MX480 additionally had a 4x40G blade and a 16x10G blade for use in connecting the local servers and networking infrastructure. AGLT2 connected its Dell/Force10 S4810 to the MX480 with 2 40G links and 8 10G links. Routing was enabled on the S4810 which allowed all of the production systems at AGLT2 to connect to the showroom floor. In addition a new storage system based upon a Dell R720 headnode, two MD3260 and two MD3060e storage shelves were directly connected to the MX480 via two 40G connections. The headnode had 2xE5-2665 processors, 256GB of RAM, a Mellanox dual ported 40G NIC (MCX314A-BCBT) and 4 SAS HBAs which connected to the 240 disks on the storage shelves. For SC12, 4 of AGLT2's production dCache servers were configured to participate in the demos (each server had 2x10G connections into the S4810).

OpenFlow and Multipath TCP for Data Intensive Science

While full use of 100 Gbps links was being shown on dedicated links, the Caltech team at SC12 also used software defined networks and multipath protocols to overcome the limitations faced by most data intensive science projects, who use shared network infrastructures or dedicated links of 10 Gbps or lower speeds.

Researchers from iCAIR, SARA, SURFnet, and Caltech demonstrated the utilization of MultiPath TCP (MPTCP) and OpenFlow to address these needs. A dynamic multipath switching fabric based on the Floodlight OpenFlow controller and Pronto OpenFlow switches in Salt Lake City, Chicago, Amsterdam and Geneva was used to create a set of link-disjoint paths across an intercontinental network. Corresponding OpenFlow forwarding entries were automatically pushed to the OpenFlow switches, whenever a new flow appeared at the ingress of the OpenFlow network. In addition, MPTCP was used to split a TCP data stream into parallel TCP sub-streams from source to destination. Using multiple paths simultaneously, an aggregate data transfer rate was reached that exceeded the maximum capacity of a single link, namely 14 Gbps on two 10 Gbps paths between Geneva and Amsterdam. The team also filled four disjoint paths from Geneva to the show floor in Salt Lake City with capacities of 5 Gbps to 10 Gbps each.

The next round of tests will target state of the art throughput across shared R&E networks, with dynamic reconfiguration of the network using Openflow triggered either by degradation of a network segment, or shifting loads from competing flows.

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

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, notably the Natural Sciences and Engineering Research Council of Canada, through the following grants: US LHCNet (DOE DE-FG02-08-ER41559), the NSF-funded ANSE (NSF-OCI 1246133) and PLaNetS (NSF PHY-0622423) projects, the DOE-funded OliMPS project (DE-SC0007346), and the US LHC Operations Program funded jointly by DOE and NSF (DOE DE-AC02-07CH11359).

We'd like to thank our industry partners involved in this year's demonstrations for providing equipment and support: ADVA, Alcatel Lucent, Ciena, Cisco, Data Direct Networks, Dell Force 10, FusionIO, IBM, Juniper, Level3, Mellanox, OCZ, and Supermicro; and our network partners BCnet, CANARIE, CENIC, Internet2, NLR, PacWave, SCinet, Starlight and US LHCNet.

Quotes on the significance of the demonstrations:

"This year the HEP team was able to demonstrate the full capability and maturity of the current generation of 100 Gbps networks and 40 GE servers, when combined with FDT using TCP, or RDMA over Ethernet. Another important demonstration was the use of OpenFlow and MultiPathTCP to build a dynamic software-defined switching fabric and send a single large stream over several disjoint paths from source to destination, in the US, Europe and across the Atlantic By sharing our methods and tools with scientists in many fields, we aim to further enable the next round of scientific discoveries, taking full advantage of 100 Gbps networks now, and higher speed networks in the near future. In particular, we hope that these developments will afford physicists and students throughout the world the opportunity to participate directly in the LHC's next round of discoveries as they emerge."

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

"At SC12 we showed that it is possible to move data reliably and efficiently between two distant sites at speeds of 100 Gbps. Collectively we were able to move data at speeds in excess of 300 Gbps and within a few years it may be possible to reach 1 Terabit/second. These achievements are critical to projects such as ATLAS on the LHC, which includes physicists from Canada. The network makes it possible to exploit the computational resources in a country that is distant from CERN and enables the researchers to play aleading role in the experiment."

-- *Randall Sobie, University of Victoria Professor and Institute of Particle Physics of Canada Research Scientist*

"This is an important step in the evolution of LHC infrastructure. The demonstrations at SC12 have shown that suitably configured sites can make effective use of 100G

infrastructure **now.** With the LHC moving into Long Shutdown 1 (LS1) at the beginning of 2013, we have until 2015 to update our global infrastructure to be ready to support the next level of discovery possible with the upgraded LHC. The work done for SC12 will be important for guiding and enabling the appropriate capabilities for the increased workload anticipated when the LHC restarts after LS1"

--- Shawn McKee, University of Michigan Research Scientist, Director of the ATLAS Great Lakes Tier-2 Center and ATLAS Physicist

About Caltech:

With an outstanding faculty that has been honored with 32 Nobel prizes and 66 National Medals of Science and Technology, 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 and a premier institution of learning. The Institute conducts instruction in science and engineering for a student body of approximately 950 undergraduates and 1,400 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 has 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 University of Victoria:

The University of Victoria, one of Canada's leading universities, provides both students and faculty with a unique learning environment. UVic has earned a reputation for commitment to research, scholarship and co-op education. The university is widely recognized for its innovative and responsive programs, interdisciplinary and international initiatives, and a diverse and welcoming learning community. For more information, see http://www.uvic.ca .

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 59,933 on all campuses. The main campus is in Ann Arbor, Michigan, and has 42,716 students (fall 2011). With over 600 degree programs and $1.2 billion in 2010-2011 research expenditures, the university is one of the leaders in innovation and research. For more information, see http://vpcomm.umich.edu/forum/michigan.php

About Brookhaven National Lab:

Established in 1947 on Long Island, Upton, New York, Brookhaven is a multi-program national laboratory operated by Brookhaven Science Associates for the U.S. Department of Energy (DOE). Seven Nobel Prizes have been awarded for discoveries made at the Lab. Brookhaven has a staff of approximately 3,000 scientists, engineers, technicians and support staff and over 4,000 guest researchers annually. Brookhaven National Laboratory's role for the DOE is to produce excellent science and advanced technology with the cooperation, support, and appropriate involvement of our scientific and local communities.

The fundamental elements of the Laboratory's role in support of the four DOE strategic missions are the following: 1) to conceive, design, construct, and operate complex, leading edge, user-oriented facilities in response to the needs of the DOE and the international community of users, 2) to carry out basic and applied research in long-term, high-risk programs at the frontier of science, 3) To develop advanced technologies that address national needs and to transfer them to other organizations and to the commercial sector, and 4) to disseminate technical knowledge, to educate new generations of scientists and engineers, to maintain technical capabilities in the nation's workforce, and to encourage scientific awareness in the general public.

About 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.

About BCNET:

BCNET is a not-for-profit, shared information technology services consortium led by its members, British Columbia’s research universities and institutes. Owned, governed and funded primarily by its members, BCNET facilitates a unique, collaborative inter-

institutional environment—one based on common goals—to explore and evaluate shared IT solutions for mutual technology challenges. BCNET's shared services approach is designed to help member institutions cut costs, minimize technology duplication, improve operational efficiency and better serve their students, faculty, staff and communities.

About CANARIE:

CANARIE Inc. is Canada's Advanced Research and Innovation Network. Established in 1993, CANARIE manages an ultra high-speed network that supports leading-edge research and big science across Canada and around the world. One million researchers, scientists and students at over 1,100 Canadian institutions, including universities, colleges, research institutes, hospitals, and government laboratories have access to the CANARIE Network. Together with 12 provincial and territorial advanced network partners, CANARIE enables researchers to share and analyze massive amounts of data, like climate models, satellite images, and DNA sequences that can lead to groundbreaking scientific discoveries. CANARIE is a non-profit corporation supported by membership fees, with the major investment in its programs and activities provided by the Government of Canada.

CANARIE keeps Canada at the forefront of digital research and innovation, fundamental to a vibrant digital economy. For additional information, please visit: www.canarie.ca.

About SARA:

In 1971 SARA was founded by the University of Amsterdam, VU University Amsterdam and CWI, the national research center for mathematics and computer science in the Netherlands. It was especially dedicated to high performance computing. Nowadays, SARA is an independent organization which supports scientific research by offering the most advanced services and expertise in the areas of computing, data storage, visualization, networking, cloud and e-Science. In 2010 SARA accommodated its ICT services for the commercial market at daughter company Vancis. SARA now exclusively addresses to the support of scientific research. For more information, please visit: https://www.sara.nl/about-sara

About USLHCNet:

US LHCNet is a Wide Area Network (WAN) infrastructure project that supports the high energy physics research mission of the U.S. Department of Energy and the National Science Foundation, with a focus on the U.S. Large Hadron Collider (LHC) program. US LHCNet supports the agencies' and the high energy physics community's strategic and tactical goals and objectives by providing interoperable, effective, and reliable transatlantic communications infrastructure and leading-edge network facilities and services in support of the US LHC program, as well as the other major programs supported by the DOE Office of High Energy Physics (OHEP).

The highest priority mission of US LHCNet is to provide reliable, dedicated, high

bandwidth connectivity between the U.S. Tier1 centers and CERN (e.g., the LHC Optical Private Network (OPN) physical links). An additional mission is to support high bandwidth traffic flows as needed, between the U.S. LHC Tier1 and European Tier2 centers as well as between the U.S. Tier2 centers and European Tier1 centers. For more information, please visit: http://www.uslhcnet.org

About iCAIR

The International Center for Advanced Internet Research (iCAIR) has been established to provide a focal point for leading edge Internet research, innovation, and early deployment. iCAIR undertakes basic R&D research projects, designs and implements leading edge prototypes, and manages several large scale advanced communication facilities. An international research center, the majority of iCAIR projects are established with its world-wide partners. For more information, please visit: http://www.icair.org/about/index.html

About Internet2

Internet2 is an advanced networking consortium led by the research and education community. An exceptional partnership spanning U.S. and international institutions who are leaders in the worlds of research, academia, industry and government, the Internet2 community is developing breakthrough network technologies that support the most exacting applications of today—and spark the most essential innovations of tomorrow.

Led by its members and focused on their current and future networking needs since 1996, Internet2 blends its unsurpassed human, IP and optical networks to develop and deploy revolutionary Internet technologies. Activating the same partnerships that produced today's Internet, our community is actively promoting the kinds of collaboration and innovation that can have—and has already had—a fundamental impact on the future of the Internet, and its fast-growing world of users. For more information, please visit: http://www.internet2.edu

About Padtec S/A:

Padtec is based in Brazil and is the largest optical communications manufacturer in South America, providing commercial equipment and solutions to over 30 countries in all continents. It has a strong focus on technological innovation and is a leader in providing solutions for large telecommunication networks, including long haul, metropolitan, access, and storage applications. Padtec partners with important telecom operators and public and private communication companies, and has distinguished itself through its presence in the trunking networks of the largest telecommunication service providers of Latin America. It also has a long tradition of successful partnerships with research networks, providing solutions to several Latin American R&E networks, such as the RedClara network in Mexico, the RedeComep, GIGA, and KyaTera networks in Brazil, as well as to ANSP (the state of São Paulo R&E network) and to RNP (the Brazilian NREN). Padtec is based in Campinas-São Paulo, Brasil, and has offices in Argentina, Peru, Mexico and Israel. For more information visit: www.padtec.com.br.

About the National Science Foundation:

The National Science Foundation (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 $6.9 billion (FY 2010), we are 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:

The U.S. Department of Energy Office of Science is the lead federal agency supporting fundamental scientific research for energy and the Nation's largest supporter of basic research in the physical sciences. The Office of Science portfolio has two principal thrusts: direct support of scientific research and direct support of the development, 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 24,000 researchers from universities, other government agencies, and private industry each year.

Contacts:

Harvey B. Newman, Caltech Professor of Physics, newman@hep.caltech.edu
(626) 395-6656

Contact: Kathy Svitil, Caltech Media Relations, ksvitil@caltech.edu
(626) 395-8022