

For Release, Sunday Nov 23 2014
**Intelligent Software Driven Dynamic Hybrid Networks With Terabit/sec Science
Data Flows**

***Caltech Network Team within the Caltech HEP: Intelligent
Software Driven Dynamic Hybrid Networks With Terabit/sec Science
Data Flows***

***New Orleans, Supercomputing 2014— Intelligent Software Driven
Dynamic Hybrid Networks With Terabit/sec Science Data Flows —***

During Supercomputing 2014 (SC14) in New Orleans, Caltech Network Team within the Caltech HEP performed a set of state of the art demonstrations entitled Intelligent Software Driven Dynamic Hybrid Networks With Terabit/sec Science Data Flows. “These tests signify the emergence of a new networking paradigm where the management of multilayer terascale networks is driven by semi-autonomous intelligent software systems, operating on an unprecedented scale with a new level of efficiency and control.” Said Professor Harvey Newman of Caltech who leads the team. “We use the Large Hadron Collider physics program to prototype these tests: it is the source of the biggest scientific data today and the physics requirements keep stretching network and storage architectures,” said Professor Maria Spiropulu at the Caltech booth on the SC14 exhibit floor.

Demonstration Setup

The on-floor SC14 Terascale Network formed an optical triangle among the Caltech, the International Center for Advanced Internet Research (iCAIR) and Vanderbilt University booths. It was also connected locally through a dark fiber to the University of Michigan booth. The layer 1 (optical) network was provided by Padtec through the use of three reconfigurable optical add-drop multiplexers (ROADMs) driven with software developed by a team from UNICAMP (Campinas), and the layer 2 (switching) network was provided by Brocade and Extreme Networks through the use of OpenFlow-capable network switches with many 100G and 40G connections, and a Software Defined Network (SDN) controller developed at Caltech, based on the OpenDaylight framework. The high density processing and storage systems supporting the highest throughput included Echostreams servers equipped with processors from Intel, 40GE network interfaces from Mellanox, and solid state disks from Seagate and Intel, as well as Dell Servers equipped with FusionIO solid state storage, and systems from SGI.

The monitoring and control of the system at both Layer 1 and 2 was coordinated through the use of Caltech's MonALISA system, and the very high

throughput was achieved through the use of Caltech's Fast Data Transfer open source TCP application.

The setup at the Caltech booth with more than 1 Terabit/sec of on-floor capacity and four 100G wide area connections, sustained 1.5 Terabits/sec of throughput during an initial memory-to-memory trial and reached its main goal of 1.0 Terabits/sec between storage and memory, including up to 400 Gigabit/sec over the wide area networks linking the conference in New Orleans to Caltech, CERN, Victoria, Michigan, and San Paulo, on the final day of the exhibit. The latest generation servers and networks deployed at SC14 and over the wide area represented a prototype of a state of the art global-scale autonomous network system, and a faithful representation of the global operations and cooperation that is inherent in global science programs handling hundreds of petabytes per year, such as the high energy physics experiments at the LHC. "Our setup encapsulates a new concept of Consistent Integrated Operations that reaps the benefits of major network development, network tools and systems, and integrates these with the mainstream data operations of the science program" said Professor Newman. The Compact Muon Solenoid (CMS) experiment has started to use dynamic circuits to manage worldwide data transfers, and the ATLAS experiment similarly makes strategic use of the network in distributed processing and data analysis.

“This is a set of groundbreaking technical achievements that we have been working towards for some time now -- the capabilities for the LHC program, and scientific computing are dramatically enhanced” said Artur Barczyk, Senior Science Network Researcher of the Caltech Network Team within the Caltech HEP

In 2015 the team will further integrate network awareness and strategic network use to support data analysis by thousands of physicists who access data remotely with local caching, generating a large set of small flows with an entirely different traffic pattern in time. Building on its close relationship with the former Center for Advanced Computing Research in 1996 - 2014, the team will from now on work closely with Caltech’s recently launched Center for Data Driven Discovery <http://cd3.caltech.edu/index.html> on strategic data operations, and associated big data analytics.

From the network point of view, the software defined network (SDN) controller operations at this year's demonstrations mark a watershed in terms of coordinated operations among the network providers and the major science users. By limiting the aggregate set of flow allocations to a "high water mark", monitoring the flows and capacity in real-time, and varying the mark as needed to adapt to changing conditions, the science programs and network providers will be able to operate the network infrastructures at high throughput levels without saturation, and avoid the destruction of competing traffic.

The methods, tools and systems we are demonstrating this year represent a major step on the way towards meeting the challenges of the future machine-to-machine communication-dominated world envisaged in the Internet of Things, with billions of data sources and sinks, and the multi-Petabyte to Exabyte scale data operations of the future High Luminosity LHC and the Square Kilometer Array, to mention two of the biggest science data projects planned in the next decade.

“We are looking forward to continuing the ambitious development trajectory and meet the future challenges with the science, network, and corporate partners” said Professor Sergio Novaes of Sao Paulo State University, who collaborated with the Caltech group in the SC14 network challenge with his team.

The Caltech and Partner Teams

Founded in 1984, and working in support of the LHC program since 1994, the Caltech Network Team within the Caltech HEP group is a worldwide leader in scientific network development, production, and operations. The Caltech team is collaborating with university teams from Michigan, UT Arlington, Vanderbilt, Victoria in Canada, UNICAMP and the State University of Sao Paulo in Brazil, and laboratory groups engaged in network development from the DOE’s Lawrence Berkeley Lab, Fermilab, and Brookhaven National Lab. The team is working with many network partners as well, including DOE’s ESnet, Internet2, CENIC, Florida

Lambda Rail, MiLR and other leading US regional networks, BCNET in Canada, leading exchange points including Starlight, AmLight, NetherLight, and CERNLight, along with GEANT, SURFNet and other European research and education networks, as well as the RNP national network and the ANSP (Sao Paulo) regional network in Brazil, on novel network system development and optimization projects focused on LHC and related applications for the last 15+ years.

This year's demonstrations were made possible in part by support from the U.S. National Science Foundation Directorate for Computing and Information Science, the U.S. Department of Energy Offices of High Energy Physics and Advanced Scientific Computing, Cisco Research, and the funding agencies of the international partners in Canada and Brazil.

The SC2014 demonstrations also were made possible through Century Link's provisioning of multiple wide area 100G links to SC2014, including a dedicated 100G link between CENIC and the Caltech campus provided by Wilcon, and the new transatlantic 200G ring provided by the ANA-200 consortium. Rapid setup of the networks was facilitated through test and measurement equipment from Spirent.

For more information, press only:

Caltech Media Relations

Deborah Williams-Hedges

Senior Media Relations Representative

(626) 395-3227

debwms@caltech.edu

For more information on Intelligent Software Driven Dynamic Hybrid

Networks With Terabit/sec Science Data Flows:

Team Lead and Contact: Harvey Newman, (626)-395-6656, newman@hep.caltech.edu

Monitoring results: See <http://sc-repo.uslhcnnet.org/display>

Government Laboratory, Network & Technology Industry Partners

ESnet	www.es.net
Fermilab	www.fnal.gov
Brookhaven National Lab	www.bnl.gov
Lawrence Berkeley Nat'l Lab	www.lbl.gov
CERN	www.cern.ch
Internet2	www.internet2.edu
CENIC	www.cenic.org
SURFNet	www.surf.nl/en/about-surf/subsidiaries/surfnet
BCNET	www.bc.net
Padtec	www.padtec.com.br
RNP	www.rnp.br
ANSP	www.ansp.br
Extreme Networks	www.extremenetworks.com
Brocade Networks	www.brocade.com
CenturyLink	www.centurylink.com

Wilcon	www.wilcon.com
Echostreams	www.echostreams.com
Intel	www.intel.com
Mellanox	www.mellanox.com
Spirent	www.spirent.com
Starlight	www.startap.net/starlight/
Manlan	noc.manlan.internet2.edu/
AmLight	www.amlight.net
CERNLight	cernlight.web.cern.ch/cernlight/