



## **SCinet Research Sandbox Shows Off Groundbreaking Network Research**

SEATTLE, Wash. – November 14, 2011 - This week, SC11 will not just showcase the next generation of HPC applications but it will also be home to eleven of the most innovative network research projects through a special program called the SCinet Research Sandbox. SCinet is the primary network infrastructure built each year for SC exhibitors to show off their most cutting edge computing applications and collaborations. As a key component of SCinet, the SRS program is designed to allow researchers to experimentally test and demonstrate their ideas on innovative network architectures, applications and protocols in the unique live environment of the SCinet network. This year, the SRS will provide researchers with access to 100 Gigabits per second dedicated links as well as a 10 Gigabit per second (Gbps), multi-vendor OpenFlow network testbed.

“In addition to supporting the extreme demands of the HPC-based demonstrations that have become the trademark of the conference, SCinet also seeks to foster and highlight developments in network research,” said Brian Tierney, SRS co-chair for SC11. “Both 100 Gbps networking and OpenFlow are poised to become some of the most influential networking technology in this decade. SRS allows the community to showcase these innovations in their infancy to demonstrate the impact they may have on the entire HPC community in the very near future.”

SRS is this year working in partnership with the Technical Program’s Disruptive Technologies (DT) track. Complementary to the mission of the SRS, Disruptive Technologies, which has taken place as part of the SC technical program since 2006, examines new computing and networking architectures and interfaces that will significantly impact the high-performance computing field throughout the next five to 15 years, but have not yet emerged in current systems.

“The Disruptive Technology program at SC is aimed at showcasing technologies and innovations that have the potential to transform high-performance computing. OpenFlow’s ability to ‘program the network’ stands to be one of the most disruptive innovations HPC, an arena that depends on, and pushes, the capabilities of network infrastructure,” said Martin Swamy, co-chair of the SRS, chair of Disruptive Technologies, and associate professor at Indiana University.

Eleven projects have been selected as part of the SRS program and the top six will showcase as part of a Disruptive Technologies program panel. For detailed information on the projects and their presentations visit: <http://sc11.supercomputing.org/?pg=scinetsandbox.html>

### **SRS projects include:**

*Advanced Individualized OpenFlow Networks With Extensions for Automated Multi-domain Topology Discovery and Provisioning Over International Hybrid Environments*

Presented by: NCHC, NCKU, National Kaohsiung University Of Applied Sciences, ICAIR at Northwestern University, iGENI Consortium, SARA High Performance Computing and e-Science Center, Communications Research Centre

Demonstrated in booths: 313, 2615, 642

This demonstration shows an automatic network topology discovery mechanism based on multi-controller OpenFlow networks. We have extended OpenFlow to include additional capabilities to enable federated domains, and to allow individualized and hybrid networks. This will be shown on an international network research testbed.

*The Data Superconductor: Demonstrating 100Gb WAN Lustre Storage Using OpenFlow Enabled Traffic Engineering*

Presented by: Indiana University with collaborators from Brocade, Ciena, Data Direct Networks, IBM, Internet2, Whamcloud, and ZIH

Booth 2239

This demonstration will feature a compute cluster and Lustre storage system taking full advantage of the 100Gbps wide area network link and OpenFlow. The collaborators will demonstrate how to empower scientists to collaborate across geographically distributed administrative domains using multi-site workflows and distribute their data from instruments to compute resources.

*Demonstration of High Performance 100 Gbps Networking for Petascale Science*

Presented by: NASA, MAX, International Center for Advanced Research (iCAIR) at Northwestern University, Laboratory for Advanced Computing at University of Chicago, Open Cloud Consortium, Ciena, Alcatel, StarLight, MREN, Fujitsu, Brocade, Force 10, Juniper, Arista, ESnet, Internet2

Booths 615, 2615, 635

These demonstrations are extensions of a series of projects that are advancing the development of 100G networks within the US and internationally, by developing high performance services necessary for the next generation of data intensive science research and discovery. Enabling large-scale time-efficient data flows over wide areas is a key issue for many advanced research disciplines.

*Enabling Large Scale Science using Inter-domain Circuits over OpenFlow*

Presented by: Internet2 and USC ISI

Booth 1327

The objective of this project is to facilitate collaboration between research groups interested in the emerging paradigms of network virtualization and software defined networking. The collaborators are developing systems which integrate the control of technologies like OpenFlow with dynamic provisioning across existing research and education infrastructures to enable research affinity groups to spontaneously form.

*End-to-End Virtualization – Campus, WAN and Data Center*

Presented by: Lawrence Berkeley National Lab, Indiana University, University of Delaware

Booth 512

This demonstration showcases the ability to build "zero configuration circuits" leveraging emerging OpenFlow technology in combination with virtual circuit technology. As part of the demo, eXtensible Session Protocol (XSP) allows GridFTP to use Remote Direct Memory Access (RDMA) enabling the application to directly interact with the transport layer. By stitching together these disparate network virtualization and flow management techniques, the demonstration showcases the power of these combined technologies to create an easily configurable end to end network path.

*FlowScale: Distributed Load Balancing of Network Traffic Using an OpenFlow Switch*

Presented by: Indiana University

Booth 2239

This is a demonstration of a new technology called FlowScale which is a traffic LoadBalancer as a service application built on OpenFlow. It is designed to provide a simple interface to load balance traffic and act as a building block to build scalable high capacity infrastructure.

*Global Scale 40Gbps RDMA Demo*

Presented by: Orange Silicon Valley, InfiniBand Trade Association and OpenFabrics Alliance

Booth 6010

This is first time transcontinental demonstration at 40Gbps of Remote Direct Memory Access (RDMA) protocols simultaneously using InfiniBand extension and iWARP. The demonstration illustrates three usage models that benefit from RDMA at such a distance: the ability to move multiple simultaneous

uncompressed video streams; the ability to access a remote data set as though it were locally mounted for visualization; and the ability to access “big data” sets at high bandwidths with low latency. The RDMA protocols in the wide area enable the three usage models to achieve up to 96 percent utilization of the available network bandwidth, improve the application processing efficiency from approximately 20 percent up to 80-90 percent and increase the messaging rate.

*HPC Application Network Performance Tuning for Global-scale Networks*

Presented by: NTT, NICT, GEMnet

Booth 4717

NTT laboratories will conduct novel HPC applications' network performance tuning over global-scale networks. Different techniques, for example, delay- and jitter- tolerance modification, traffic-characteristics optimization and fine-grained network load balancing, are performed in a domain-by-domain manner based on our globally distributed high-precision network measurements systems (PRESTA10G).

*Monitoring large Ethernet networks with Ethernet OAM enabled OpenFlow controllers*

Presented by: SARA High Performance Computing and e-Science Center, iCAIR at Northwestern University, Laboratory for Advanced Computing University of Chicago, Communications Research Centre  
Booth 642

This demonstration showcases how OpenFlow can be used by end-users to easily add new network protocols to OpenFlow switches. Protocols, in this case a protocol to detect link failures, are implemented in an OpenFlow Controller once and can then be used with any switch that supports the OpenFlow API to enable the protocol on that switch.

*RouteFlow: Virtualized IP Routing Services in OpenFlow networks*

Presented by: CPqD

Booth 2239

The main goal of RouteFlow is to develop an open-source framework for virtual IP routing solutions over commodity hardware implementing the OpenFlow API. This demonstration will showcase how RouteFlow can work within the commodity routing architecture that combines the line-rate performance of commercial hardware with the flexibility of open source routing stacks running on general purpose computers.

*Steroid OpenFlow Service (SOS): A Transparent Network Performance Enhancement Service*

Presented by: Clemson University

Booth 5401

This demonstration intends to show a novel paradigm for network operators to seamlessly introduce performance enhancement and other network services. With an OpenFlow controller that detects specific flow types (e.g., TCP in this demo) and redirects them to operator provided service agents (e.g., parallel TCP service), SOS demonstrates how operators can leverage OpenFlow to dispatch customized services to users in a campus network or across a wide area network.

**About SC11**

SC11, sponsored by the ACM (Association for Computing Machinery) and the IEEE Computer Society, offers a complete technical education program and exhibition to showcase the many ways high performance computing, networking, storage and analysis lead to advances in scientific discovery, research, education and commerce. This premier international conference includes a globally attended technical program, workshops, tutorials, a world class exhibit area, demonstrations and opportunities for hands-on learning. For more information on SC11, please visit: <http://sc11.supercomputing.org/>.

###