

November 07, 2016 | By Jan Zverina

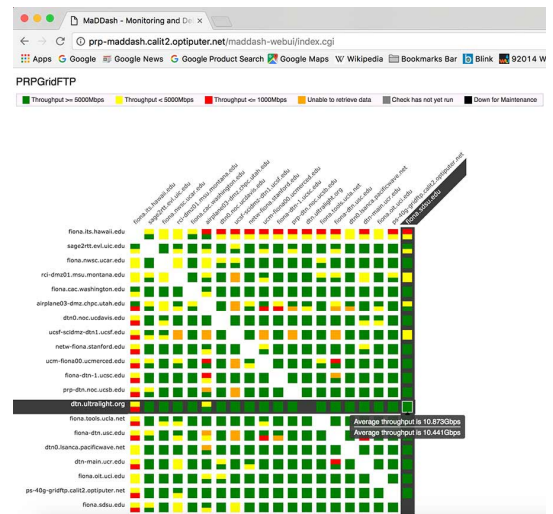
SDSC to Host High-Speed, Large Data Transfer Experiment at SC16 Show

Goal is to push limits of data transfer in UC San Diego’s Pacific Research Platform Project

Scientific discovery requires investigations of extremely large amounts of data, which must be transported at a high level of performance among multiple distributed collaborator sites. The Pacific Research Platform (PRP) initiative is addressing the many challenges of such large-scale transport by implementing network architecture and technology not common on general networks.

The San Diego Supercomputer Center (SDSC) at the University of California San Diego will host a node in the PRP at this year’s International Conference for High Performance Computing, Networking, Storage and Analysis (SC16), in Salt Lake City, Utah November 14-17. The aim is to achieve large data transfers over long-distance networks to the limits possible using single personal computers of the type that researchers may have in their labs in the near future.

This specially-designed data transfer node (DTN) was developed by the Department of Energy’s Energy Science Network (ESnet), enhanced with the newest solid-state “disk” drives (called “NVMe” for Non-Volatile Memory express). Each NVMe allows terabytes (TBs) of disk-to-disk transfer at speeds approaching the networks’ maximum rate of 100 gigabits/second. The SDSC node at SC16 is one of four so-called FIONAs (for Flash Input/Output Network Appliances) built by the PRP for this experiment. The FIONAs will be

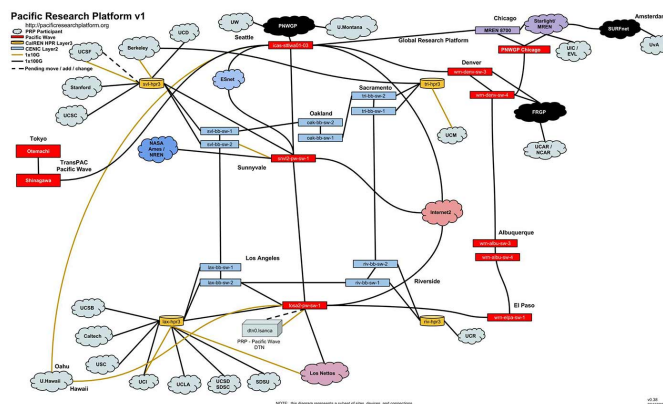


Twenty PRP partner institutions transfer 10GB of data disk-to-disk four times daily using GridFTP and ESnet’s MaDDash visualization software. Most transfers exceed 5Gb/s (green) and some exceed 10Gb/s, as shown between Caltech and SDSU, two universities featured at SC16 trying to approach 100Gb/s disk-to-disk. The PRP is funded by NSF Award #ACI-1541349, and also made possible by support from CENIC and the University of California Office of the President.

located at Caltech and San Diego State University, and the SDSC and Caltech booths at SC16. The four nodes will be joined by others at Northwestern University and at the SC16 StarLight booth, among others, at the annual conference.

In 2015, the National Science Foundation announced funding of a \$5 million, five-year award for UC San Diego and UC Berkeley to establish the PRP, a high-capacity, data-centric “freeway system” that will eventually give participating universities and other research institutions the ability to move data about 1,000 times faster than speeds on today’s inter-campus shared internet.

“Our goal is to demonstrate that open-source data transfer protocols, such as fast data transfer (FDT) from Caltech, can move data over thousands of miles on best-effort networks, approaching a rate that moves a terabyte – or one million, million bytes – disk-to-disk in less than 12 minutes, and sustain this extraordinary rate for one hour, until the 6TB disks fill,” said Tom DeFanti, a research scientist at the California Institute for Telecommunications and Information Technology (Calit2)’s Qualcomm Institute, located on the UC San Diego campus.



Pushing the limits of data transfer exposes limitations in the networks and the network switches, as well as the operating systems and data transfer software sending and ingesting the data. The PRP routinely moves data disk-to-disk among its approximately 20 partners at speeds over 10 gigabits/second.” More details are at <http://prp-maddash.calit2.optiputer.net/maddash-webui/index.cgi>.

“The goal of these experiments at SC16 is to run much faster so we fail, debug the networks and DTNs end-to-end, and ultimately succeed and share the information with the SC community,” said DeFanti. “These kind of experiments during the past 20-plus years have built the research-supporting networks by first proving them at SCInet and in booths at the annual SC shows.”

Visualizations of the successes and failures of the data transfer will be shown in the SDSC booth in real time by John Graham, PRP’s senior development engineer. Additional experiments will test Software Defined Networking (SDN) technology between the SDSC booth and UC San Diego’s MicroCloud server. Several international high-speed data transfers on new

links will also be attempted. Caltech Senior Network Engineer Azher Mughal has been the primary architect and organizer of these links between California and the SC16 venue in Salt Lake City.

Following the conclusion of SC16, many of the PRP partner FIONAs will simultaneously send and receive data to participate in stress tests of the networks as directed by SCInet engineers.

Larry Smarr, founding director of Calit2 and Professor of Computer Science and Engineering at UC San Diego, is Principal Investigator of the PRP. He and his staff wish to thank SDSC for hosting the demonstrations, as well as Caltech, SDSC and UC San Diego, the Corporation for Education Networking in California (CENIC), Pacific Wave, StarLight, Century Link, Wilcon, and SCInet for their leadership in 100Gb/s networking.

MEDIA CONTACT

Jan Zverina, 858-534-5111, jzverina@sdsc.edu

UC San Diego's [Studio Ten 300](#) offers radio and television connections for media interviews with our faculty, which can be coordinated via studio@ucsd.edu. To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.