

Researchers building high-speed data 'freeway system'

August 5 2015



Pacific Research Platform principal investigator Larry Smarr looks at Calit2's Vroom display system. The new data network will allow partner institutions to experiment with real-time sharing between ultra-resolution display systems.

For the last three years, the National Science Foundation (NSF) has made a series of competitive grants to over 100 U.S. universities to

aggressively upgrade their campus network capacity for greatly enhanced science data access. NSF is now building on that distributed investment by funding a \$5 million, five-year award to UC San Diego and UC Berkeley to establish a Pacific Research Platform (PRP), a science-driven high-capacity data-centric "freeway system" on a large regional scale. Within a few years, the PRP will give participating universities and other research institutions the ability to move data 1,000 times faster compared to speeds on today's inter-campus shared Internet.

The PRP's data sharing architecture, with end-to-end 10-100 gigabits per second (Gb/s) connections, will enable region-wide virtual co-location of data with computing resources and enhanced security options. PRP links most of the research universities on the West Coast (the 10 University of California campuses, San Diego State University, Caltech, USC, Stanford, University of Washington) via the Corporation for Education Network Initiatives in California (CENIC)/Pacific Wave's 100G infrastructure. To demonstrate extensibility PRP also connects the University of Hawaii System, Montana State University, the University of Illinois at Chicago, Northwestern, and the University of Amsterdam. Other research institutions in the PRP include Lawrence Berkeley National Laboratory (LBNL) and four national supercomputer centers (SDSC-UC San Diego, NERSC-LBNL, NAS-NASA Ames, and NCAR). In addition, the PRP will interconnect with the NSF-funded Chameleon NSFCloud research testbed and the Chicago StarLight/MREN community.

"Research in data-intensive fields is increasingly multi-investigator and multi-institutional, depending on ever more rapid access to ultra-large heterogeneous and widely distributed datasets," said UC San Diego Chancellor Pradeep K. Khosla. "The Pacific Research Platform will make it possible for PRP researchers to transfer large datasets to where they work from their collaborators' labs or from remote [data centers](#)."

Multicampus collaboration

Fifteen existing multi-campus data-intensive application teams act as drivers of the PRP, providing feedback over the five years to the technical design staff. These application areas include accelerator particle physics, astronomical telescope survey data, gravitational wave detector data analysis, galaxy formation and evolution, cancer genomics, human and microbiome 'omics integration, biomolecular structure modeling, natural disaster, climate, CO2 sequestration simulations, as well as scalable visualization, virtual reality, and ultra-resolution video. The PRP will be extensible both across other data-rich research domains as well as to other national and international networks, potentially leading to a national and eventually global data-intensive research cyberinfrastructure.

"To accelerate the rate of scientific discovery, researchers must get the data they need, where they need it, and when they need it," said UC San Diego computer [science](#) and engineering professor Larry Smarr, principal investigator of the PRP and director of the California Institute for Telecommunications and Information Technology (Calit2). "This requires a high-performance data freeway system in which we use optical lightpaths to connect data generators and users of that data."

The leadership team includes faculty from two of the multi-campus Gray Davis Institutes of Science and Innovation created by the state of California in the year 2000: Calit2, and the Center for Information Technology Research in the Interest of Society (CITRIS), led by UC Berkeley.

"The Pacific Research Platform is an ideal vehicle for collaboration between CITRIS and Calit2 given the growing importance of universities working together for the benefit of society," said CITRIS Deputy Director Camille Crittenden, co-PI on the PRP award. "The project also

received strong support from members of the UC Information Technology Leadership Council, which includes chief information officers [CIOs] from the 10 UC campuses, five medical schools, the Lawrence Berkeley National Lab and the Office of the President."

Crittenden will manage the science engagement team and the enabling relationships with CIOs on participating campuses and labs.

With all 10 UC campuses involved, the UC Office of the President was so convinced of the PRP's value that it provided additional funding to "maintain the momentum" until NSF funds could become available. "This cyberinfrastructure evolution shows the transformative power of uniting all ten campuses and providing a platform for the ubiquitous flow of knowledge for research collaboration," said Tom Andriola, the UC System Chief Information Officer. The leading IT administrators of the non-UC institutions are also strongly committed to the PRP.

The PRP is basing its initial deployment on a proven and scalable network design model for optimizing science data transfers developed by the U.S. Department of Energy (DOE) – the ESnet Science DMZ. "ESnet developed the Science DMZ concept to help address common network performance problems encountered at [research institutions](#) by creating a network architecture designed for high-performance applications, where the data science network is distinct from the commodity shared Internet," said ESnet Director Greg Bell, a division director at Lawrence Berkeley National Lab (LBNL). "As part of its extensive national and international outreach, ESnet is committed to working closely with the Pacific Research Platform to leverage the Science DMZ and Science Engagement concepts to enable collaborating scientists to advance their research." In the PRP the Science DMZ model will be extended from a set of heterogeneous campus-level DMZs to an interoperable regional model.

As fast as local disks

"PRP will enable researchers to use standard tools to move data to and from their labs and their collaborators' sites, supercomputer centers and data repositories distant from their campus IT infrastructure, at speeds comparable to accessing local disks," said co-PI Thomas A. DeFanti, a research scientist in Calit2's Qualcomm Institute at UC San Diego.

DeFanti and co-PI Phil Papadopoulos, a program director in the San Diego Supercomputer Center (SDSC), will coordinate the efforts of the large group of network engineers, network providers and measurement programmers at the PRP institutions.

The PRP project emerged from earlier NSF grants awarded to the PRP investigators (OptIPuter, GreenLight, StarLight, Quartzite, and Prism), which led to brainstorming at the CENIC 2014 annual retreat. A subsequent one-day workshop in December 2014, hosted at Stanford, led to a decision to publicly demonstrate the feasibility of the PRP. To do so, the partners engaged network engineers from a number of PRP member campuses to work intensively for the first 10 weeks of 2015 on a proof-of-principle demonstration of high-performance data transfers between Science DMZs over existing elements of the proposed infrastructure. This required extensive collaboration among the PRP partner campuses, led by CENIC's John Hess. The resulting PRPv0 was presented at CENIC 2015 on March 9, 2015 at UC Irvine. The result involved disk-to-disk data transfers from within one campus Science DMZ to another. These transfers, using 100GbE infrastructure, are made possible because of the NSF CHERub award that upgraded UC San Diego to the CENIC bandwidth.

After iteration and tuning, tests demonstrated data transfer speeds of 9.6Gb/s out of 10 from UCB, UCI, UCD, and UCSC to UC San Diego, with two transfers at 36Gb/s out of 40 from UCLA & Caltech to UC San

Diego. During the demonstration at CENIC 2015, one PRP-optimized test moved 1.6 Terabytes in four minutes; by contrast, using the default campus Internet, it took three hours to transfer 0.1 Terabytes, demonstrating a 720x improvement.

These experiments, and the opportunity to create the PRP itself, were only possible because of the CENIC's leadership in connecting every research university in California at 100G, while upgrading its backbone to 100G. This \$10M+ investment provided a catalyst and leverage for individual campuses to upgrade their Science DMZ infrastructure, using either local funds or NSF resources. "The CENIC demonstration in March showed the commitment and expertise of the partnership members," recalled CENIC President Louis Fox. "Based on that carefully monitored demonstration, we are convinced that the enhanced [infrastructure](#) of the PRP can succeed in its ambitions."

Separately, NSF has awarded funds to hold a PRP design workshop at UC San Diego, now scheduled for October 2015, entitled: 'Building an Interoperable Regional Science DMZ.' This workshop will bring together the PRP application driver researchers with the distributed computer architects, the network engineers, and the multi-institutional IT/Telecom administrators to further refine the PRP implementation.

Science focus

In addition to DeFanti, Papadopoulos, and Crittenden, Frank Würthwein, a physicist at UC San Diego and SDSC program director, is a PRP co-PI; he will lead technical development of the application groups and monitor progress from the scientists' perspective.

"The PRP is not a build-it-and-they-will-come exercise," said Würthwein, who is also executive director of the Open Science Grid. "The cyberinfrastructure is responsive to the existing and expected needs

of data-intensive applications, so we are building a very science-focused platform that will put these universities above and beyond what other regions already have."

Würthwein is closely involved in the global Large Hadron Collider (LHC) community, which accounted for roughly two-thirds of Open Science Grid's 800 million computational hours in 2014. Other disciplines consuming OSG resources include social sciences (notably economics), engineering and medicine. The PRP-wide LHC cyberinfrastructure is a direct outgrowth of the SDSC LHC UC-wide initiative, started in October 2014.

The PRP will be rolled out in two phases. First, the PRPv1 platform will focus on deploying its data-sharing architecture to include all member campuses. Once all of the institutions are up and running, the consortium will develop and then offer PRPv2 as an advanced, IPv6-based version with robust security and software-defined networking (SDN) features.

The computers that "terminate" the optical fiber Big Data flows in DMZ systems, sending, receiving, measuring, and monitoring data, are termed by ESnet Data Transfer Nodes (DTNs). Within each campus Science DMZ, the Pacific Research Platform will deploy a DTN developed at UC San Diego under the NSF-funded Prism@UCSD project, led by PRP co-PI Papadopoulos. Dubbed Flash I/O Network Appliances (FIONA), they are modestly-priced, Linux-based computers made of commodity parts, featuring terabytes of flash drives optimized for data-centric applications.

"FIONAs act as data super-capacitors for the Science Teams," said Papadopoulos. "They already serve a similar purpose in the Prism@UCSD project that interconnects two-dozen big-data laboratories on the UC San Diego campus with XSEDE resources at SDSC. Prism@UCSD is the university's Science DMZ, with a core switch

router whose lit bandwidth is now approaching 1 Terabyte per second." The Prism@UCSD network allows connected labs to use burst bandwidth cross-campus that would oversaturate and likely impair the campus backbone, which serves 50,000 users daily.

UC San Diego Computer Science and Engineering Chair Rajesh Gupta said, "We are proud that a member of our department faculty, Larry Smarr, is once again providing visionary leadership on a large-scale project that will have a transformative impact on national cyberinfrastructure. PRP envisions a practical distributed architecture supporting a wide range of disciplines to ensure that federally funded university research advances science and continues to produce extraordinary talent for generations to come."

Provided by University of California - San Diego

Citation: Researchers building high-speed data 'freeway system' (2015, August 5) retrieved 18 June 2024 from <https://phys.org/news/2015-08-high-speed-freeway.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.