

October 16, 2013 | By Jan Zverina

CHERuB Connection to Accelerate Data-Driven Science at UC San Diego

NSF-funded “information superhighway” to connect UC San Diego with 100gbps data networks

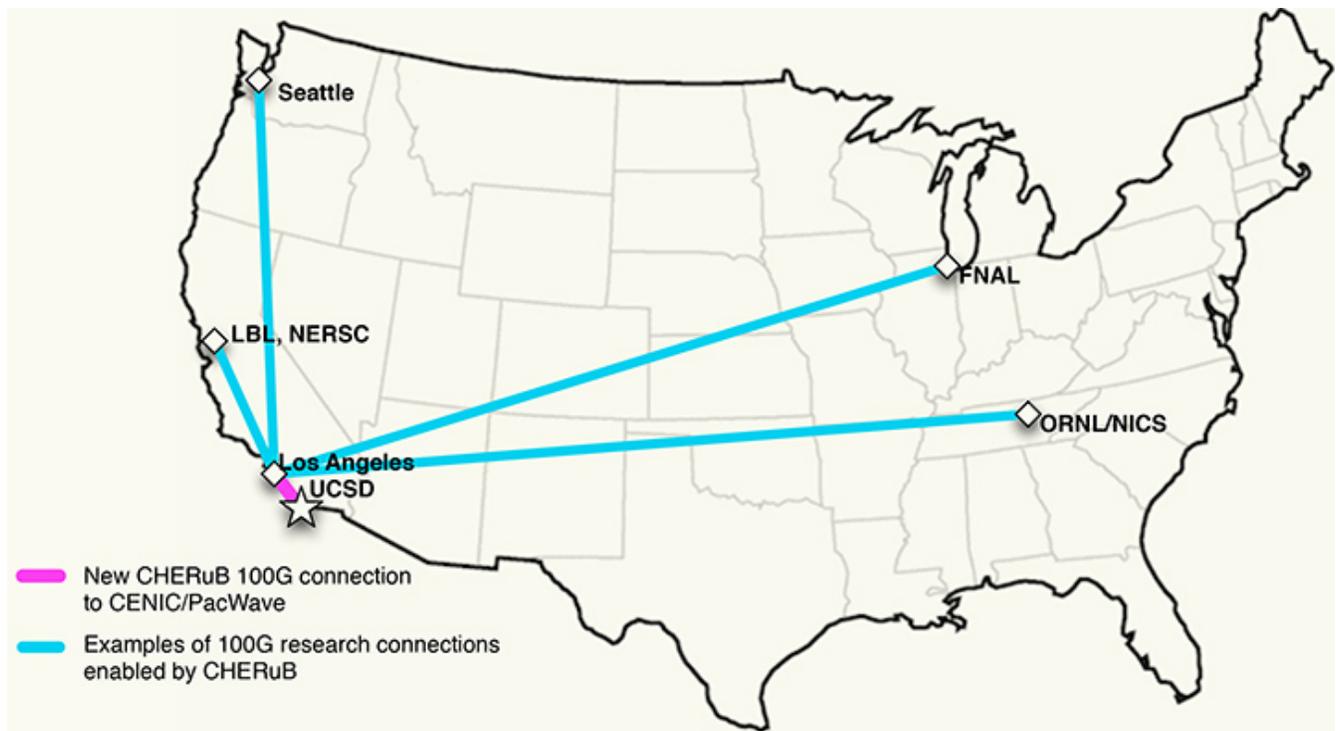


Image showing new CHERuB 100 Gigabit-per second connection to CENIC/Pacific Wave, and some examples of other 100Gbps connections enabled by CHERuB to key national labs such as Fermilab (FNAL) and Oak Ridge National Laboratory/National Institute for Computational Science (ORNL/NICS). Image by Valerie Polichar, UC San Diego

The San Diego Supercomputer Center (SDSC) at the University of California, San Diego, and the university’s Administrative Computing and Telecommunications (ACT) organization have been awarded a National Science Foundation (NSF) grant to connect the campus to high-bandwidth national research networks to help advance a new range of data-driven research.

Named CHERuB for Configurable, High-speed, Extensible Research Bandwidth, the project is funded under a two-year, \$500,000 award from the NSF's Division of Advanced Cyberinfrastructure, starting January 1, 2014. The initiative will provide 100Gbps (Gigabits per second) connectivity – the new high-end for wide-area research networks – to support multi-institutional data transit over networks such as the [Internet2's Advanced Layer 2 Service \(AL2S\) \(PDF\)](#) and [ESnet](#), as well as a joint project between those networks called the [Advanced Networking Initiative \(ANI\)](#), the result of a \$62 million grant under the American Recovery and Reinvestment Act to build a national 100G “information backbone.”

When completed, the CHERuB link will place UC San Diego among research universities and institutions having the highest available connectivity, with a capacity 10 times greater than existing modern data networks.

“This final piece of connectivity puts UC San Diego on the information superhighway of the future,” said Sandra A. Brown, Vice Chancellor for Research at UC San Diego. “CHERuB will make possible a wide range of new instructional activities that cross institutions and locations, and is likely to benefit other rapidly emerging areas of science that increasingly rely on big data exchanges while removing the network itself as a bottleneck to scientific discovery.”

Added Min Yao, the university's Assistant Vice Chancellor of Administrative Computing and Telecommunications: “CHERuB will make UC San Diego's already-impressive production network into a true research network. Multiple scientists will be able to simultaneously send large data streams to and from the Internet, allowing research advances across a variety of disciplines.”

Examples of research domains that will benefit from CHERuB include large cosmology, atmospheric sciences, electron microscopy, genomic sequencing, oceanography, high-energy physics, and telemedicine – all of which can encompass data-rich research and rely on multi-site or inter-institutional activities.

The CHERuB award comes on the heels of an earlier NSF award of the same amount to build a research-defined, end-to-end cyberinfrastructure on UC San Diego's campus capable of supporting enormous bursts of data between facilities. That project, called [Prism@UCSD](#), was announced in March.

“CHERuB is the missing piece that will connect UC San Diego's Prism network to even faster national networks to advance scientific research,” said SDSC Director Michael Norman, principal investigator for the project. “One might compare it to if the Interstate-5 freeway

running alongside UC San Diego had no exits to the university. Consider this project as building the on-ramps and off-ramps to the campus.”

Nobel Physics Prize Connection

UC San Diego’s 100G connectivity to ESNet means that the campus will be connected to all of the Department of Energy’s (DOE) Office of Science laboratories and supercomputer centers, including [Fermilab \(FNAL\)](#), the Tier-1 data center for the U.S. portion of the [Compact Muon Solenoid \(CMS\)](#) project that is involved in research that led to this year’s Nobel Prize in Physics. SDSC’s *Gordon* supercomputer has been providing auxiliary computing capacity by processing massive data sets generated by the U.S. Compact Muon Solenoid, one of two large general-purpose particle detectors at the Large Hadron Collider used by researchers to find the elusive [Higgs particle](#).

“The combination of CHERuB and Prism@UCSD is a game-changer for us,” said Frank Wuerthwein, a professor of physics at UC San Diego and a member of the CMS project. “In combination with the technologies developed by the NSF-funded [“Any Data, Anytime, Anywhere”](#) project, compute resources in San Diego can now realistically have direct access to disk resources at FNAL and vice versa. This significantly simplifies and at the same time accelerates getting our science done.”

Several other projects now underway will benefit directly from the increased bandwidth, serving as test beds for the deployment, measurement, and ongoing monitoring of the CHERuB network. Specifically, CHERuB will:

- Enable scientists throughout the country to transmit very large data sets to SDSC to take advantage of the data analysis and visualization services of *Gordon* as well as [Comet](#), a new supercomputer to be deployed by SDSC in early 2015;
- Provide a means for users to submit [Open Science Grid](#) jobs in an efficient and effective manner;
- Allow the use of the [Dynamic Network System \(DYNES\)](#), a nationwide cyber-instrument spanning about 40 U.S. universities and 11 Internet2 connectors, as well as other virtual scientific instruments, over much larger pipes;
- Provide a high-bandwidth external path to [SDSC’s Cloud Storage and Archive infrastructure](#);
- Enable the building of on-demand high-bandwidth “superhighways” supporting both native IPv4 and IPv6, from UC San Diego researchers to important national research networks and to other educational sites in California;

- Develop best practices guides and other resources for researchers using new networks that can be shared with other campuses making the move to high bandwidth connections; and
- Continue to provide a cutting edge IT infrastructure for UC San Diego research and instruction.

“This exciting project will make it easier for NSF- and DOE-funded scientists across the United States to benefit from SDSC’s data-intensive *Gordon* supercomputer, and eventually *Comet*, said Gregory Bell, director of ESnet and director of the Scientific Networking Division at Lawrence Berkeley National Laboratory. “We look forward to working with CENIC, Internet2, UC San Diego and SDSC to make data move even faster for high energy physics, genomics, and many other fields.”

CENIC View

The project also includes an upgrade of the campus gateway to support CENIC’s newly-installed 100Gbps pip between San Diego and Pacific Wave’s 100Gbps regional research network collective in Los Angeles, and intra-campus infrastructure to facilitate high-speed access by targeted research activities.

“CENIC is committed to working closely with UC San Diego to make this 100Gbps connection a reality,” said Louis Fox, CENIC’s president and CEO. Added Internet2 President and CEO H. David Lambert: “Internet2 has a long history of collaboration with SDSC and UC San Diego. Among other things, this connection will allow the project to facilitate high-speed communication by the Open Science Grid with various Advanced Networking Initiative sites.”

The CHERuB project is funded under NSF grant # ACI-1340964 under the provisions of NSF 13-530. Co-principal investigators of the project are Valerie E. Polichar, infrastructure architect at ACT; and Thomas E. Hutton, SDSC’s network architect and manager.

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>.