

CTWatch QUARTERLY [May 2005](#)

Introduction

GUEST EDITORS

Philip Papadopoulos, Director, Advanced Cyberinfrastructure Lab, San Diego Supercomputer Center
Larry Smarr, Harry E. Gruber Professor, Department of Computer Science and Engineering, Jacobs School of Engineering and Director, California Institute for Telecommunications and Information Technology, University of California San Diego

Welcome to the second issue of the Cyberinfrastructure Technology Watch Quarterly. In this issue we focus on the state of one and 10 Gbps long-haul, optical circuits supporting the research community. It has been over a decade (1994) since the Very High-Speed Backbone Network Services (vBNS) connected major NSF centers and universities at OC-3 (155 Megabits/Sec Mbps). Soon thereafter, selected circuits were upgraded to OC-12 (622Mbps). In 1997, Internet2¹ was formed to connect a larger collection of universities at OC-12 and Gigabit speeds. Internet2 now operates with 10 Gigabit backbones. However, according to sites such as NASA's ENSIGHT,² end-to-end file transport from major scientific data repositories to end users laboratories across the shared internet is less than 1% of that: typically 50-100 mbps.

In the late 1990s, the future looked promising for long-haul research networks, and in three years vBNS/Internet2 dramatically increased long-haul research network capacity. But in 2000, the "Tech Bubble" turned into the "Tech Meltdown." During the Bubble phase, fiber-strands were being laid worldwide at the rate of over 8,000 km/hour (70 million kilometers in 1999)! Multiplying this build-out, Dense Wave-Division Multiplexing (DWDM) enabled carriers to run multiple 10-gigabit or even 40-gigabit channels (each termed a "lambda" for the wavelength bin in the infrared band it occupies) on a single fiber pair. This meant that the carriers had a functional way to multiply bandwidth without expensive trenching of new fiber. Current DWDM technology supports up to eighty 10-gigabit "waves" or "lambdas" on such networks (a mere 800 times the capacity of Internet2 backbone of eight years ago).

As a number of long-haul network companies went bankrupt, a new opportunity became available because of the bandwidth overbuild—telecom carriers were willing to discuss long term leases of fiber or lambdas to individuals. Foreign countries took the lead, with Canada's CANARIE being in the vanguard. In the United States, NCSA, Argonne National Lab, EVL at UIC, and Northwestern convinced the state of Illinois in 1999 to extend the Illinois Century Network to construct a dark fiber state network, called I-WIRE,³ to support Illinois researchers' needs for large amounts of bandwidth. In 2001, the Distributed TeraGrid Facility (DTF) proposed connecting large data nodes and computing clusters on a national scale using a 40 Gigabit dedicated backbone (four 10 Gb lambdas) among the four originating centers to form what we know now as the TeraGrid.⁴ Also in 2001, NSF funded the international "point of entry" for research networks STARTAP to become StarLight⁵ – a 1GigE and 10GigE switch/router facility for high-performance access to participating networks and a true optical switching facility for wavelengths.

We see these three state, national, and international initiatives as the catalyzing events. Researchers worldwide were convinced that large, dedicated optical circuit research networks were not only theoretically possible but were practically being put into service. A scant four years after the original TeraGrid award and three years after I-WIRE became operational, there are now eight sites connected to the extended TeraGrid, using the new formed National LambdaRail (NLR)⁶ to create extensions of the original four lambdas, and by now over two dozen state and regional dark fiber networks exist and are interconnecting to NLR.

In this issue, we are fortunate to have three articles whose authors have all played critical roles in this new age of long-haul research networks. They delve deeper into the details and provide critical insights:

Linda Winkler from Argonne National Laboratory has been in the trenches for both Teragrid and SciNet (the SCxy Conference's big monster network that exists for 5 days every November). In her article, "Does 10G Measure Up?", she describes the challenges of high-end research deployments and highlights heroic efforts in the Bandwidth Challenge where current state-of-the-art clocks in at over 100Gbps for an application using multiple 10G networks at SC2004. Linda also takes us through some of the Teragrid infrastructure.

In an article titled "The National LambdaRail" by Dave Farber and Tom West, the authors describe how regional research networks have taken advantage of the abundance of dark fiber to enable multiple, research-focused 10 Gigabit networks. Farber and West give a very nice condensed history of high-speed networking, how a variety of environmental factors made the NLR possible, how the NLR is being used today and what researchers might expect looking out 10 years.

Fast Research Networks are not a US-only concession. In fact, some would say that the US is a fast follower in the on-demand, lambda-based network. In "Translight, A Major US Component of GLIF", Tom DeFanti, Maxine Brown, Joe Mambretti, John Silvester, and Ron Johnson describe the optical interconnections available between U.S. and international researchers. Truly, research networks have gone global and big/fast research networks are becoming prevalent. International partners are critical in the world of Team Science.

Here at UCSD, we are building a campus-level OptIPuter⁷ interconnecting five laboratories and clusters of three functional types: compute, storage, and scalable tiled display walls. The total number of nodes in the OptIPuter fabric exceeds 500 and each lab has four fiber pairs that connect it to a central high-speed core switching complex that has both a Chiaro Enstara⁸ (a large router based on a unique optical core) and standard Cisco⁹ 6509 switch-router. The current instantiation supports both 10 gigabit and one gigabit signals running from each lab to the central core. In a follow-on NSF-funded proposal, Quartzite augments this structure by adding DWDM signaling on the established fiber plant, a transparent optical switch from Glimmerglass, and in 2006, a wavelength selective switch from Lucent. When complete, the Quartzite switching complex will be able to switch packets, wavelengths or entire fiber paths, allowing us to build different types of network layouts and capabilities to test OptIPuter research and other optical networking ideas. With reconfigurable networks and clusters, OptIPuter/Quartzite forms a campus-scale research instrument. At build-out, this instrument will support nearly half a Terabit of lambdas landing into a central, reconfigurable complex.

OptIPuter and Quartzite preview what campuses need to evolve to: immense bandwidth, optical circuits on demand, and reconfigurable endpoint systems. Of critical importance is the evolution of large and network-capable storage clusters that can be accessed with clear paths from research labs scattered around campus. Using cluster management systems (we use the Rocks Clustering Toolkit¹⁰), most scalable systems (compute clusters, tiled display clusters, application servers) can be thought of as soft-state. However, as science moves to the inevitable data-intensive modes, information storage is critical to the campus research enterprise. This coming generation of campus networks allows storage silos (critical state) to be remote from labs, then managed and operated on behalf of researchers without losing performance or adaptability for the research scientists themselves. In essence, soft-state systems can be put anywhere on campus (notably in labs), and critical-state systems are not required to be in physical proximity. "Unlimited" campus network capacity allows universities to co-optimize the preservation of critical data and the ability to rapidly change soft-state systems to meet research challenges.

We'd like to close with the following thought. Long haul, fast research networks are springing up everywhere and bandwidth is finally meeting the "it will be abundant" predictions that many of us have believed for nearly a decade. However, the missing link overall is the campus connectivity – some campuses are pioneering big networks, but most still operate on one gigabit backbones. It is a strange turn of events when the long-haul network is fatter and more capable than your campus network.

Arden Bement, the director of the National Science Foundation recently discussed this issue in the Chronicle of Higher Education.¹¹

“Research is being stalled by ‘information overload,’ Mr. Bement said, because data from digital instruments are piling up far faster than researchers can study. In particular, he said, campus networks need to be improved. High-speed data lines crossing the nation are the equivalent of six-lane superhighways, he said. But networks at colleges and universities are not so capable. “Those massive conduits are reduced to two-lane roads at most college and university campuses,” he said. Improving cyberinfrastructure, he said, “will transform the capabilities of campus-based scientists.”

¹ <http://www.internet2.org/>

² http://ensight.eos.nasa.gov/active_net_measure.html

³ <http://www.iwire.org/>

⁴ <http://www.teragrid.org/>

⁵ <http://www.startap.net/starlight/>

⁶ <http://www.nlr.net/>

⁷ <http://www.optiputer.net/>

⁸ <http://www.chiaro.com/>

⁹ <http://www.cisco.com/>

¹⁰ <http://www.rocksclusters.org/>

¹¹ Kiernan, V. “NSF Has Plan to Improve ‘Cyberinfrastructure,’ but Agency’s Director’s Gives Few Details,” The Chronicle of Higher Education, Volume 51 (36), May 2005.

<http://chronicle.com/prm/weekly/v51/i36/36a03001.htm>

URL to article: <http://www.ctwatch.org/quarterly/articles/2005/05/introduction/>