NASA GSFC Input for OptIPuter's FY06 Report

Prepared by J. Patrick Gary, 6/21/06

As an OptIPuter Partner, during March 2006 NASA GSFC assisted the Calit2/OptIPuter/CAMERA (Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis) Project in stress testing a new 10-Gbps lambda-based connection obtained from the NLR between Chicago, IL and McLean, VA by generating 9.2 gigabit per second (Gbps) of User Datagram Protocol (UDP)-based Internet Protocol (IP) packets between two pairs of GSFC-based computers, but where the data traffic between those pairs was looped out to Chicago, IL, and back (http://cisto.gsfc.nasa.gov/L-Netpdfs/Opti-L_test_config_032106.pdf). These test filled the NLR-STAR-WASH-10GE-103 10-GigE circuit to 92% of capacity in both directions simultaneously (http://cisto.gsfc.nasa.gov/L-Netpdfs/Opti-10g_032106.pdf). These tests were run for 6 hours, which each resulted in over 22.68 TB of data being transmitted simultaneously in each direction. In several 6-hour duration tests of the new connection, the NLR never lost a packet, and the Calit2/OptIPuter/CAMERA Project accepted the link for operational use before April 1, 2006. At Mclean GSFC further interconnected that NLR-based 10-Gbps link with a 10-Gbps port on a DRAGON switch/router which then enabled a 1-Gbps link across the DRAGON regional optical network to the J. Craig Venter Institute in Rockville, MD.

On April 3, 2006, in collaboration with Dr. Robert Grossman (UIC) and the TeraFlow Testbed (TFT) project, GSFC connected a TFT-provided cluster of four dual-Opteron 265's servers to GSFC's Force10 E600 10-GE switch at GSFC. Each TFT server is 10-GE connected via an Intel Pro1000 NIC to a TFT-provided SMC SMC8708L2 10-GE switch which connects via 10-GE with the Force10 E600 at GSFC. The TFT cluster at GSFC subsequently has been used to conduct "teraflows" with other TFT clusters in Chicago, Illinois; Kingston, Ontario; Amsterdam, The Netherlands; Geneva, Switerland; Tokyo, Japan; and London, England. The TFT cluster at GSFC uses the L-Net's 10-Gbps pathway across GSFC's channel49 lambda on the DRAGON and WASH-STAR lambda on the NLR to connect with the OptIPuter's Force10 E1200 10-GE switch at the StarLight facility in Chicago. TFT's network diagram (http://www.ncdm.uic.edu/maps/index.jpeg) shows GSFC's connection in the TFT. GSFC is hoping later to leverage TFT's connection into Tokyo to enable high performance access to Coordinated Enhanced Observing Period (CEOP) project data sets hosted at the University of Tokyo.

Also based on experience and information gained as an OptIPuter Partner, NASA GSFC further implemented the following NASA 10-Gbps network connections during the FY06 period of performance.

Recognizing the demand for reliable high-speed networks to support NASA's

high performance computing community and its projected data flow requirement increases, GSFC's local-area Scientific and Engineering Network (SEN) and the GSFC-based NASA Center for Computational Sciences (NCCS) supercomputer facility each upgraded their network infrastructures to 10 Gbps. These upgrades capitalize on the ARC-managed wide-area NASA Research and Engineering Network (NREN) upgrade, which increased the bandwidth in its backbone link between GSFC and ARC's Project Columbia supercomputer from 1 Gbps to 10 Gbps using the National LambdaRail (NLR).

On March 3, 2006, the SEN's perimeter router (a Force10 E600 10-Gigabit Ethernet (10-GE) switch/router) began operationally supporting a 10-Gbps connection from the NREN's GSFC-local 10-GE switch/router (a Cisco 6506). In the weeks prior, GSFC's Bill Fink assisted the NREN in stress testing their new 10-Gbps lambda-based connection obtained from the NLR by generating 9.2-Gbps of UDP-based IP packets between two pairs of SEN-based computers (http://cisto.gsfc.nasa.gov/L-Netpdfs/L-Net_mrtg_NREN-L_020306.pdf), but where the data traffic between those pairs was looped out to Sunnyvale, CA, and back (http://cisto.gsfc.nasa.gov/L-Netpdfs/NREN-L_test_config_020106.pdf). In several 6-hour duration tests of the new connection, the NLR never lost a packet, and the NREN Project accepted the link for operational use.

On April 19, 2006, the NCCS began pre-operational end-to-end readiness tests of its 10-Gbps network infrastructure, which includes a near-10-Gbps firewall capability designed and implemented by the NCCS's Lee Sheridan (CSC), multiple 10-GE ports on its new Force10 E600 10-GE switch/router, and 10-GE network interface cards (NICs) for several NCCS supercomputer platforms. NCCS' firewall capability is essentially an Intel-based Linux box, with Neterion 10-GE NICs in PCI-X2 I/O buses, running IP Tables stateful rulesets. Bill Fink's stress testing that new firewall capability demonstrated that the NCCS's new firewall capability supports up to 6.7-Gbps of network throughput performance. Fink's stress testing was accomplished using locally-generated single-stream Transport Control Protocol (TCP)-based IP packets between a pair high-performance workstations with Myricom 10-GE NICs in their PCI-Express I/O buses. End-to-end single-stream TCP-based IP packet flow testing between an NCCS SGI Origin 3800 and a not-fully-tuned high-performance workstation at ARC across the 10-Gbps NREN demonstrated up to 1.5-Gbps throughput performance, and testing with two simultaneous TCP-based streams demonstrated up to 3.0-Gbps throughput performance.

Operational cut-over of the entire set of NCCS supercomputer platforms to their 10-GE network infrastructure was accomplished on June 7, 2006; and the SEN expects to deploy a 10-Gbps firewall capability for the protection of other SEN users before the end of July 2006.

In recent presentations made to the GSFC Sciences and Exploration Directorate's Data Archive and Distribution for High-Performance Computing Working Group, Christa Peters-Lidard, Head of GSFC's Hydrological Sciences Branch, concluded with the prediction, "NASA science will be bandwidth limited – not CPU limited." Mike Seablom, Head of GSFC's Software Integration and Visualization Office, added, "Now and beyond, data and networking are the central elements to enable faster leaps in performance." The 10-Gbps upgrades of the NREN, SEN, and NCCS network infrastructures will help fulfill those predictions.