

ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS)

Chin Guok and David Robertson

Energy Sciences Network, Lawrence Berkeley National Laboratory

Summary

Two DOE Office of Science workshops in the past two years have clearly identified dynamically provisioned quality of service (QoS) paths as a science discipline driven network requirement. To address this, the ESnet On-Demand Secure Circuits and Advance Reservation System (OSCARS) is developing and deploying a prototype service that enables on-demand provisioning of guaranteed bandwidth, secure virtual circuits within the ESnet production network.

With the Large Hadron Collider (LHC) expected to come online in 2007, demand for guaranteed high-bandwidth connectivity for huge data transfers is becoming urgent. Such service is also identified as important by many other DOE Office of Science mission areas in the reports of several science networking requirements workshops¹, and such service will be important to the ITER project when it comes online a decade from now.

The default characteristics of the Internet today do not provide a user any service guarantees. There is neither the assurance that a packet will be delivered to its destination, nor any transport predictability (such as latency and jitter) when the packet is in transit. With the advent of service sensitive applications to support DOE's

science (such as remote controlled experiments, time constrained massive data transfers, distributed (Grid based) analysis of massive scientific data sets, etc.) the capability of on-demand, dynamically provisioned bandwidth assignment will provide predictable high bandwidth and/or low latency paths (virtual circuits) to high-impact science applications in a way that lets the science collaborations manage bandwidth as a predictable service as is currently done with computing resources.

In order to ensure that the new network services interoperate among DOE Labs and universities and other research institutions, the new services are being developed in close cooperation with Internet2/Abilene and the US research and education network community, and with Dante/GEANT and the European research and education network community. This work is also in collaboration with the DOE funded UltraScienceNet, network research project.

In the past year, ESnet has successfully demonstrated that an end-user can provision

¹ See "High Performance Network Planning Workshop," August 2002 (<http://www.doecollaboratory.org/meetings/hnpw>) and "DOE Science Networking Roadmap Meeting," June 2003 (<http://www.es.net/hypertext/welcome/pr/Roadmap/index.html>)

circuits within ESnet (through a simple web-interface), and effectively obtain the required bandwidth. We have demonstrated that as a proof-of-concept, OSCARS is a useful service that can, on a user-demand basis, provide: i. lower jitter, ii. guaranteed bandwidth and traffic isolation, and, iii. traffic-engineer flows to bypass congested links.

Lower Jitter

Tests done by Les Cottrell (*DWMI: Datagrid Wide Area Monitoring Infrastructure*) show that lower jitter can be obtained by using OSCARS circuits (Figure 1). This characteristic is vital to remote-controlled experiments and applications and video conferencing.

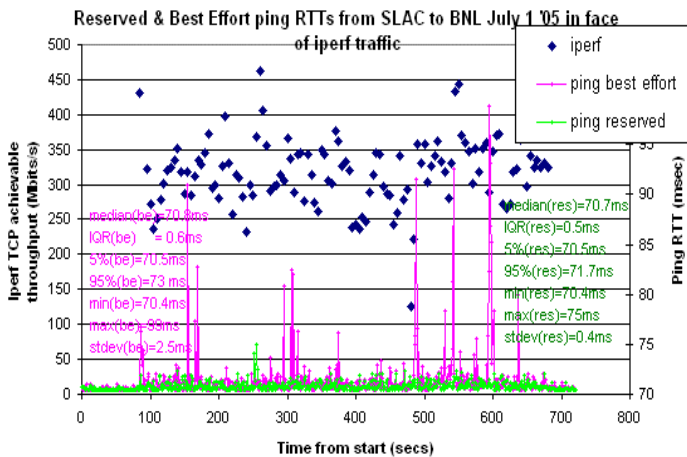


Figure 1. Ping Tests Between SLAC and BNL²

Guaranteed Bandwidth and Traffic Isolation Using OSCARS Circuits

Tests done by Dantong Yu (*TeraPaths: A QoS Enabled Collaborative Data Sharing Infrastructure for Peta-scale Computing Research*³) and Vyto Grigaliunas (*Lambda Station*⁴) show that the user-reserved bandwidth cannot be exceeded regardless of

² <http://www-iepm.slac.stanford.edu/dwmi/oscars>

³ <http://www.atlasgrid.bnl.gov/terapaths>

⁴ <http://www.lambdastation.org>

which protocol is used. This opens the possibility of running more efficient protocols that may be impolite (e.g. UDP) without adversely affecting other traffic.

Tests done by David Schissel (*Network Quality of Service for Magnetic Fusion Research*) show that the user-reserved bandwidth is effectively obtainable even when the path was completely congested by competing traffic.

Traffic-Engineer Flows to Bypass Congested Links

In the first week of April 2005, the production LHC 10GE connection between FNAL and CERN was disrupted due to a fishing boat severing CERN's link to Chicago. As a result of the cut link, the LHC Service Challenge data was rerouted over FNAL's production OC12 (622mb/s) causing severe congestion.

As a temporary measure, a persistent OSCARS circuit carrying only LHC traffic was configured between 32AoA (in New York) and Starlight (in Chicago). This caused the LHC traffic to traverse FNAL's non-production connection at Starlight, offloading the production OC12.

To date we have created 21 accounts for OSCARS beta users, collaborators, and developers, and have processed more than a hundred reservation requests.

For further information on this subject contact:

Dr. Thomas Ndousse, Program Manager
Mathematical, Information, and
Computational Sciences Division
Office of Advanced Scientific Computing
Research

Phone: 301-555-3691

tndousse@er.doe.gov