

Intradomain and Interdomain Circuit Reservations using OSCARS

Chin Guok, Evangelios Chaniotakis, Energy Sciences Network, and David Robertson*,

Summary

Two DOE Office of Science workshops⁺ in the past five 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 end-users to provision guaranteed bandwidth secure virtual circuits within the ESnet production network, and between ESnet and other domains.

The practice of modern science is increasingly dominated by large-scale science that is highly dependent on world wide collaborations of scientists and science resources. An example of this is the Large Hadron Collider (LHC) which is expected to come online in 2008. With the environment of science being inherently multi-domain, where data flows over several networks under different administrative controls, creating end-to-end guaranteed bandwidth paths poses unique challenges. To address these challenges, OSCARS was designed and engineered in collaboration with the DICE (Dante, Internet2, CANARIE, ESnet) Control Plane working group which encompasses the major R&E networks in North America and Europe, and also domestic projects such as TeraPaths¹ from Brookhaven National Laboratory, Network QoS for Magnetic Fusion Research² from General Atomics, and LambdaStation³ from Fermi National Laboratory.

Circuit Service Development

With the successful implementation and deployment of layer 3 (IP) virtual circuit (VC) services within ESnet, efforts have been re-focused on implementing a layer 2 (specifically Ethernet VLAN) VC service.

The development of the layer 2 (Ethernet VLAN) VC service was to align OSCARS with the DICE Control Plane efforts at the data plane exchange layer (i.e. multi-domain end-to-end VC will be a Ethernet VLAN service).

Collaborative Efforts

DICE

The participation of ESnet in the DICE Control Plane working group has been very productive. In the past year, the working group has reached an agreement on the topology exchange schema. This schema was authored within the framework of the OGF Network Measurements Working

¹ TeraPaths: <http://www.atlasgrid.bnl.gov/terapaths>

² Fusion QoS: <http://www.fusiongrid.org/network>

³ LambdaStation: <http://www.lambdastation.org>

* Distributed Systems Department, Lawrence Berkeley National Laboratory

⁺ See "High Performance Network Planning Workshop," August 2002

(<http://www.doecollaboratory.org/meetings/hpnpw>) and "DOE Science Networking Roadmap Meeting," June 2003 (<http://www.es.net/hypertext/welcome/pr/Roadmap/index.html>)

Group⁴ (NMWG) to promote synergy between the DICE Control Plane working group and perfSONAR⁵ (which define the framework for monitoring the VC after it is created). Dynamic topology exchange using this schema was successfully demonstrated between GEANT2, Internet2, and ESnet in Sep 2007. A protocol document outlining best practices for the topology schema implementation is currently in draft form.

Inter-domain resource scheduling, which is the next step beyond topology exchange has also been demonstrated between ESnet (OSCARS) and Internet2 (DRAGON/HOPI/DCN). The details of the interoperability exercise have been captured in an implementation document. Dynamic resource scheduling between GEANT2, Internet2, and ESnet is scheduled for SC07.

Nortel

Nortel has shown an interest in inter-operating with OSCARS and have began developing code to enable the Nortel Dynamic Resource Allocation Controller (DRAC⁶) to communicate to OSCARS using the interface as defined by OSCARS and DRAGON/HOPI/DCN. Interoperability testing is expected to commence in Nov 07.

TeraPaths and LambdaStation

With the development of a layer 2 VC service in OSCARS, additional work has been done to address the issue of interoperability between TeraPaths and LambdaStation with OSCARS and DRAGON/HOPI/DCN. (Layer 3 VC service interoperability between TeraPaths and OSCARS was established back in Aug 2006.) To facilitate this, a central repository containing pertinent documents (e.g. service

definitions) for each of the projects has been setup. An interoperability demonstration is targeted for SC07.

Universiteit van Amsterdam

The use of tokens for authentication and authorization in virtual circuit signaling has been a long-time research project of the University of Amsterdam⁷. Preliminary efforts to collaborate and integrate tokens into the OSCARS signaling phase is currently underway.

Hybrid-MLN

In collaboration with the Hybrid Multi-Layer Network Control for Emerging Cyber-Infrastructures project⁸, extensive measurements using Spirent AX4000 test gear were taken over numerous virtual circuits to determine end-to-end network properties (e.g. jitter, latency). These VCs spanned multiple domains (i.e. ESnet, DRAGON, I2, USN), and were configured using various technology capabilities (e.g. MPLS LSPs, Ethernet switched VLANs, SONET/SDH provisioned time slots, Lambda switching). The ESnet portion of the multi-domain end-to-end VC was configured as an OSCARS circuit.

The findings of the study concluded that the various network technologies performed well both individually and when concatenated together, provided that the input traffic was not driven close to or beyond the bottleneck path capacity.

For further information on this subject contact:

Chin Guok
Energy Sciences Network
chin@es.net
(510) 495-2279

⁴ OGF NMWG: <http://nmwg.internet2.edu>

⁵ perfSONAR: <http://www.perfsonar.net>

⁶ DRAC: <http://www.nortel.com/drac>

⁷ UvA SNE:

http://www.science.uva.nl/research/snet/category/04_generic_aaa/

⁸ Hybrid-MLN: <http://hybrid.east.isi.edu>