



DOE Network Research PI Meeting September 28, 2005



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

Chin Guok (chin@es.net)
ESnet Network Engineer

David Robertson (dwrobertson@lbl.gov)
DSD Computer Software Engineer

Lawrence Berkeley National Laboratory



Virtual Circuit Network Service

- Every requirements workshop involving the science community has put bandwidth-on-demand as the highest priority – e.g. for
 - Massive data transfers for collaborative analysis of experiment data
 - Real-time data analysis for remote instruments
 - Control channels for remote instruments
 - Deadline scheduling for data transfers
 - “Smooth” interconnection for complex Grid workflows New network services are critical for ESnet to meet the need of large-scale science like LHC.

What are Characteristics of Today's Flows – How “Dynamic” a Circuit?

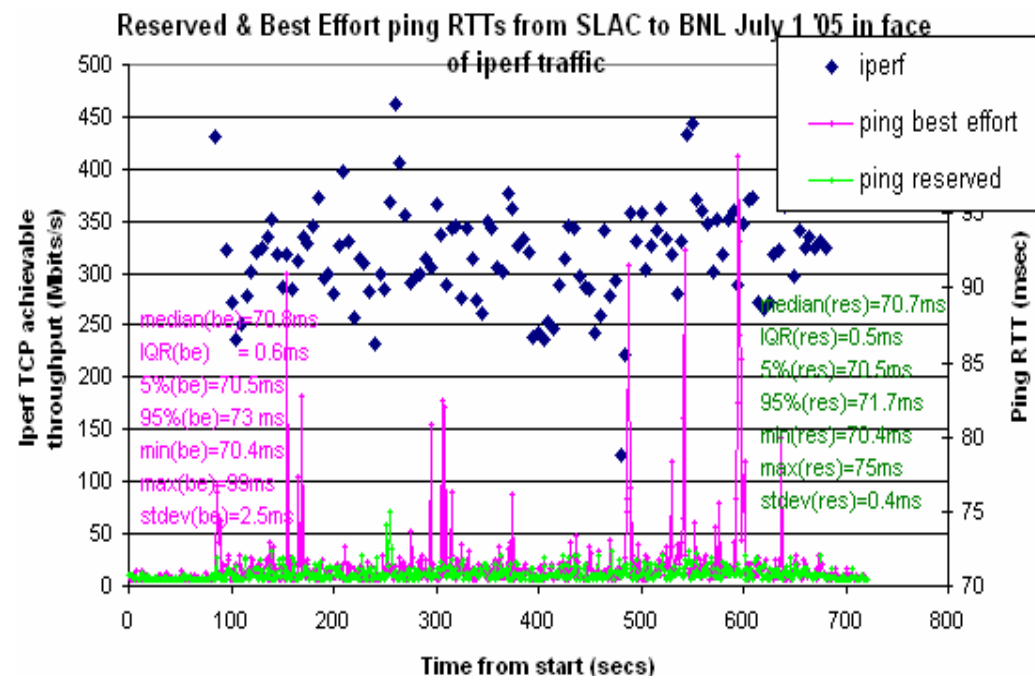
- LIGO – CalTech
 - Over 1 year the “circuit” duration is about 3 months.
- SLAC – IN2P3 (FR)
 - Over 1 year the “circuit” duration is about 1 day to 1 week.
- SLAC – INFN (IT)
 - Over 1 year the “circuit” duration is about 1 to 3 months.
- FNAL – IN2P3 (FR)
 - Over 1 year the “circuit” duration is about 2 to 3 months.
- INFN (IT) - SLAC
 - Over 1 year the “circuit” duration is about 3 weeks to 3 months.

OSCARS: Guaranteed Bandwidth Circuit Service

- New network services are critical for ESnet to meet the need of large-scale science like LHC.
- Most important new network service is ***dynamically provisioned virtual circuits*** that provide:
 - Traffic Isolation
 - Will enable the use of high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP based transport. (see, e.g., Tom Dunigan's compendium <http://www.csm.ornl.gov/~dunigan/netperf/netlinks.html>)
 - Guaranteed Bandwidth
 - The only way that we have currently to address deadline scheduling – e.g. where fixed amounts of data have to reach sites on a fixed schedule in order that the processing does not fall behind far enough so that it could never catch up – very important for experiment data analysis.

First Beta Tests with End Users (2Q05)

- DIII-D National Fusion Facility (GA)
 - Sean Flanagan of General Atomics (GA) performed several tests transferring fusion data from an MDSPlus server in GA to a client at the NERSC facility in Oakland (OSF) using OSCARS reserved circuit.
- Internet End-to-end Performance Monitoring (IEPM)
 - Les Cottrell of SLAC performed several tests with and without the set up of OSCARS reservations (<http://www-iepm.slac.stanford.edu/dwmi/oscars/>).
 - IEPM OSCARS-91
Jitter Measurement



Tough Problem: Inter-Domain Interoperability

- **Motivation:**
 - For a virtual circuit service to be successful, it must
 - Be end-to-end, potentially crossing several administrative domains
 - Have consistent network service guarantees throughout the circuit
- **Observation:**
 - Setting up an intra-domain circuit is easy compared with coordinating an inter-domain circuit
- **Issues:**
 - Cross domain authentication *and* authorization
 - A mechanism to authenticate and authorize a bandwidth on-demand (BoD) circuit request must be agreed upon in order to automate the process
 - Multi-domain Acceptable Use Policies (AUPs)
 - Domains may have very specific AUPs dictating what the BoD circuits can be used for and where they can transit/terminate
 - Domain specific service offerings
 - Domains must have way to guarantee a certain level of service for BoD circuits
 - Security concerns
 - Are there mechanisms for a domain to protect itself (e.g. RSVP filtering)

Collaborative Efforts

- Internet2/GEANT
 - ESnet hosted a 2-day working meeting with Internet2's Bandwidth Reservation for User Work (BRUW) project. (2Q05)
 - Joint development of code for OSCARS and BRUW .
 - ESnet hosted a 1-day working meeting with I2's BRUW and GEANT's JRA3 projects. (3Q05)
 - Discussion of UNI and NNI common service definitions.
 - Discussion of workflow processes for inter-domain reservations.
 - Possible integration of Authentication/Authorization (AA) work done by I2's Shibboleth and GEANT's JRA5.
 - Monthly teleconferences for status updates and technical discussions. (3Q05)
- UltraScience Net (USN)
 - Working with Nagi Rao from USN to develop a roadmap document for ESnet/USN collaboration. (3Q05)
- General Atomics (GA)
 - ESnet hosted a 1-day working meeting with GA's Network Quality of Service for Magnetic Fusion Research project. (3Q05)

Roadmap

- Testing
 - Test effectiveness of OSCARS circuits using real application (DIII-D's EFIT) with actual data in real-time. (4Q05)
 - Usage of OSCARS circuits during DIII-D experimental operations. (2Q06)
- Authentication/Authorization (AA)
 - Use Grid-Shib for OSCARS AA. (1Q06)
- Inter-domain Interoperability
 - Agree upon UNI and NNI common service definitions and workflow process with BRUW and JRA3. (2Q06)
 - Preliminary demonstration of inter-domain AA with BRUW and JRA3. (3Q06)
- Service Offerings
 - Add layer-2 VLAN circuits to OSCARS service. Currently OSCARS only offers IP circuits. (3Q06)