

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

ESCC/Internet2 Joint Techs Workshop
Jul 16-19, 2006

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

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

Neena Kaushik (nrkaushik@scu.edu)
Santa Clara University

Outline

- **Requirements for Virtual Circuit Services**
- **Virtual Circuit Characteristics**
- **Virtual Circuit Service Functional Requirements**
- **Functional View of OSCARS**
- **OSCARS Path Computation**
- **OSCARS Collaborative Efforts**
- **OSCARS: Guaranteed Bandwidth VC Service for SC Science**
- **ESnet Virtual Circuit Service Roadmap**

➤ Requirements for Virtual Circuit Services

- Identified as one of the two most important new network service by the 2002 science requirements workshop (Ref-1) (the other being end-to-end performance monitoring)
- Intended to address specific science needs

➤ Virtual Circuit Characteristics

- **Traffic isolation and traffic engineering**
 - Provides for high-performance, non-standard transport mechanisms that cannot co-exist with commodity TCP-based transport
 - Enables the engineering of explicit paths to meet specific requirements
 - e.g. bypass congested links, using lower bandwidth, lower latency paths
- **Guaranteed bandwidth [Quality of Service (QoS)]**
 - Addresses deadline scheduling
 - Where fixed amounts of data have to reach sites on a fixed schedule, so that the processing does not fall far enough behind that it could never catch up – very important for experiment data analysis
- **Reduces cost of handling high bandwidth data flows**
 - Highly capable routers are not necessary when every packet goes to the same place
 - Use lower cost (factor of 5x) switches to relatively route the packets
- **End-to-end connections are required between Labs and collaborator institutions**

➤ Virtual Circuit Service Functional Requirements

- **Support user/application VC reservation requests**
 - Information required in the request include
 - Source and destination of the VC
 - Bandwidth, start time, and duration of the VC
 - Traffic characteristics (e.g. flow specs) to identify traffic designated for the VC
- **Manage allocations of scarce, shared resources**
 - Authentication to prevent unauthorized access to this service
 - Authorization to enforce policy on reservation/provisioning
 - Gathering of usage data for accounting
- **Provide circuit setup and teardown mechanisms**
 - Widely adopted and standard protocols (such as MPLS and GMPLS) are well understood within a single domain
 - Cross domain interoperability is the subject of ongoing, collaborative development
- **Enable the claiming of reservations**
 - Traffic destined for the VC must be differentiated from “regular” traffic
- **Enforce usage limits**
 - Per VC admission control is needed to police usage, which in turn facilitates guaranteed bandwidth
 - Consistent per-hop QoS throughout the network facilitates transport predictability

➤ Functional View of OSCARS

- **Support user/application VC reservation requests**
 - ✓ Users enter reservations via a web-page
 - ✓ Applications uses API to send signed SOAP messages
- **Manage allocations of scarce, shared resources**
 - ✓ Centralized resource management
 - ✓ Authentication is done using X509 certificates
 - Authorization TBD
- **Provide circuit setup and teardown mechanisms**
 - ✓ OSPE-TE for routing
 - ✓ BGP for reachability
 - ✓ RSVP-TE for signaling
 - ✓ MPLS for switching
- **Enable the claiming of reservations**
 - ✓ Policy based filtering for traffic destined for VC
- **Enforce usage limits**
 - ✓ Per VC admission control
 - ✓ Separate router queue for VCs

➤ OSCARS Path Computation

- Extend constraint path computations to include parameters not available in RSVP-TE. (e.g. latency, jitter, peak period costing)
- Order k-path algorithm

find_path (set q)

START

s = source node; d = destination node;

set temp = q; set l; //List of paths

if (total elements(temp) == 0)

return 0;

For each entry t in temp

if (t == path from s to d)

add entry t to set l

else

subpath sp = t;

For all possible one-hop path destinations from the subpath sp

If new subpath satisfies the all constraints (e,g, bw/latency/no-cycle)

add to the set temp

Remove the subpath sp from set temp

find_path(temp);

END


➤ OSCARS Collaborative Efforts

- **To ensure compatibility, the design and implementation is done in collaboration with the other major science R&E networks and end sites**
 - Internet2: Bandwidth Reservation for User Work (BRUW) (*Ref-2*)
 - Development of common code base
 - Successful inter-domain VC reservation and setup. X.509 signed soap messages over SSL used for inter-domain communication.
 - GEANT: Bandwidth on Demand (GN2-JRA3), Performance and Allocated Capacity for End-users (SA3-PACE) and Advance Multi-domain Provisioning System (AMPS) (*Ref-3*)
Extends to NRENs
 - Instance of AMPS inter-domain manager installed in ESnet testbed.
 - Successful inter-domain reservation (no setup) between AMPS inter-domain manager at GEANT and ESnet.
 - Developing OSCARS service WSDL description to model that of the GEANT2 PACE project
 - BNL: TeraPaths - A QoS Enabled Collaborative Data Sharing Infrastructure for Peta-scale Computing Research (*Ref-4*)
 - Interoperability tests between OSCARS and Terapaths utilized WSDL description modeled from the GEANT2 PACE project
 - GA: Network Quality of Service for Magnetic Fusion Research (*Ref-5*)
 - SLAC: Internet End-to-end Performance Monitoring (IEPM) (*Ref-6*)
 - USN: Experimental Ultra-Scale Network Testbed for Large-Scale Science (*Ref-7*)
- **In its current phase this effort is being funded as a research project by the Office of Science, Mathematical, Information, and Computational Sciences (MICS) Network R&D Program**
- **A prototype service has been deployed as a proof of concept**
 - To date more than 20 accounts have been created for beta users, collaborators, and developers
 - More than 100 reservation requests have been processed

➤ **OSCARS: Guaranteed Bandwidth VC Service For SC Science**

- **ESnet On-demand Secured Circuits and Advanced Reservation System (OSCARS) (Ref-8)**
- **In its current phase this effort is being funded as a research project by the Office of Science, Mathematical, Information, and Computational Sciences (MICS) Network R&D Program**
- **A prototype service has been deployed as a proof of concept**
 - To date more than 20 accounts have been created for beta users, collaborators, and developers
 - More than 100 reservation requests have been processed

➤ ESnet Virtual Circuit Service Roadmap



Footnotes

- Ref-1 Report of the High Performance Network Planning Workshop
http://www.es.net/pub/esnet-doc/2-3high-performance_networks.pdf
- Ref-2 Internet2 BRUW Project: <http://discvenue.internet2.edu/wordpress>
- Ref-3 GEANT PACE Project: <http://pace.geant2.net>
- Ref-4 BNL TeraPaths Project: <http://www.atlasgrid.bnl.gov/terapaths>
- Ref-5 General Atomics QoS Project: <http://www.fusiongrid.org/network>
- Ref-6 SLAC IEPM Project: <http://www-iepm.slac.stanford.edu>
- Ref-7 UltraScienceNet Testbed: <http://www.usn.ornl.gov>
- Ref-8 ESnet OSCARS webpage: <http://www.es.net/oscars>