On-demand Secure Circuits and Reservation System (OSCARS)

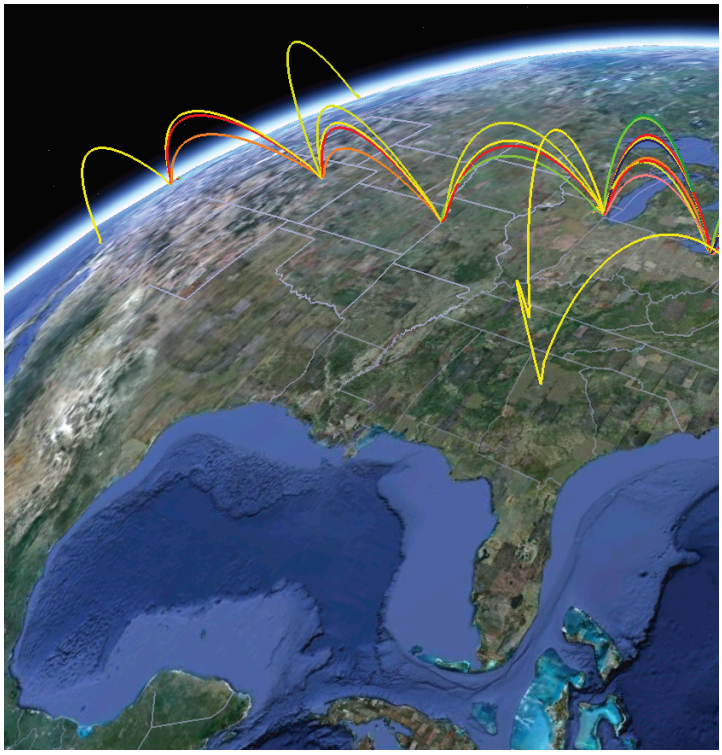

"By using OSCARS to access large data sets and learn how to operate and manage global-scale networked systems, U.S. educational institutions remain on the leading edge of research necessary to attract the best students and faculty from around the world. OSCARS is especially important to smaller sites that may not have access to crucial network and/or storage resources otherwise."

-- Michael Ernst, Director, RHIC & ATLAS Computing Facility, U.S. ATLAS Facility Manager

Overview

The increasingly data-intensive nature of next-generation research in areas including high-energy physics, climate change, genomics and other fields is challenging the capacity of the scientific networks that support universities and research institutions. Berkeley Lab's ESnet developed the On-demand Secure Circuits and Advance Reservation Software (OSCARS), which provides the ability to provision circuits between UC scientists and their collaborators that guarantee a level of performance that cannot be achieved over traditional networks. For example, researchers can remotely access scientific instruments like high-powered microscopes using a network connection that meets stringent parameters for latency. These remote connections help foster higher instrument utilization and greater scientific productivity for UC researchers.

Highlights

- *Reliable, production-quality service* — OSCARS software guarantees users network bandwidth, end-to-end throughput and quality of service in the time slots they choose, helping them to reliably access and exchange massive, time-sensitive data sets—so that competing traffic cannot cause congestion and packet loss. OSCARS is used in multiple scientific applications, including:
    - The Large Hadron Collider at CERN: 100 terabytes of data arrive daily at U.S. computing facilities, and OSCARS distributes the data to the next level of labs and research centers.
    - National Oceanic and Atmospheric Administration: historical climate data and climate models run on supercomputers, involving data exchanges with scientists around the world.
    - The U.C.'s Joint Genome Institute (JGI): Genomics data analysis and storage, where OSCARS was successfully deployed to seamlessly reroute critical data within hours of an equipment failure.

- *Automation saves time and effort* — OSCARS' web services interface lets even non-technical scientists automate the process of setting up temporary virtual circuits —accomplishing in a few minutes transactions that previously took weeks or months to do manually, with a multitude of phone calls and emails.

- *Seamless service across networks* — Like an efficient travel agent, the OSCARS interface can elegantly handle complex routing. OSCARS circuits facilitate automated transport of large data sets to their end destination over multiple connecting networks and domains.

- *Widely deployed* — OSCARS is deployed in more networks and supports more hardware platforms than any competing software. Currently OSCARS supports over 40 institutions in the U.S. and abroad, and has been adopted by R&E (research & education) networks such as Internet2 to connect California universities to the broader research and education community, CalTech to connect their particle physicists around the world, as well as international R&E networks in Brazil, Norway, and others.

- *Optimizes available bandwidth* — OSCARS enables each user to access and use more bandwidth than he or she initially reserved when available because other users are not sending traffic on their reservations. Network operators now can optimize the available bandwidth on a system,

while guaranteeing users the bandwidth they requested. In addition, OSCARS provides the flexibility to provision both short-term and long-term virtual circuits to handle data flows that can last for minutes to years, such as the one produced by the Large Hadron Collider.

- *Interoperable* — OSCARS is interoperable with other software and devices. OSCARS enabled networks can use a range of equipment from vendors including Juniper, Cisco, Ciena, Brocade, Adva/Movaz and Force10, but remain independent by building a flexible "Path Setup System" to control multi-vendor boxes. As new devices such as OpenFlow (www.openflow.org) come to market, OSCARS can be extended to support them as well.

- *Failure recovery* — OSCARS provides the flexibility to allow users to customize resiliency requirements and manage failure recovery.

- *Vehicle for Technology Transfer* — OSCARS creates a ready framework for translating innovative network research into wide deployment and production use. Since its inception, OSCARS development has been open-sourced to the greater technical and academic community. Groups at universities around the country, notably UC Davis, are developing specialized OSCARS applications for their networks that can potentially improve efficiency for the multiple R&E networks that deploy OSCARS to support scientific research around the world.

Project Details

- OSCARS was developed by ESnet PIs using funding from the DOE Office of Science (ASCR).
- The code features of OSCARS 0.5.3 were frozen in November 2010 and it was released for general availability.
- OSCARS version 0.6, offering increased modularity, was made available for download and testing by the community in 2010 at http://code.google.com/p/oscars-idc

OSCARS includes the following technical highlights:

- OSCARS uses resource scheduling and path routing based on resource availability and existing reservations, i.e., multi-constrained path computation. Both processes belong to a class of decision problems in which the time required to solve a problem increases rapidly as the size of the problem grows. OSCARS works by aggressively paring away extraneous information prior to executing the path computation.

- When users schedule bandwidth reservations via an OSCARS software interface, OSCARS authenticates and authorizes the user, then determines the optimal solution path for the user's request according to network link topology.
- The scalable path computation engine in OSCARS can determine a solution path by taking into account multiple constraints such as bandwidth (the transmission capacity of the link), latency (the time it takes to transmit a unit of data), as well as energy consumption to promote green networking, which is a growing movement to reduce the energy consumption of data centers.
- OSCARS software incorporates complex technologies that allow network operators to automate and deliver a reliable, guaranteed network service to the end-user, here described briefly:

1. **Web-browser interface/Programmatic API:** interacts directly with the user or user application to get information regarding the connectivity requirements (e.g. how much bandwidth, what is the duration, where are the end-points, etc.)

2. **Authentication/Authorization/Accounting module:** authenticates the user and authorizes the user's request. This component also provides accounting information for billing or auditing.

3. **Topology manager:** maintains and updates the network link topology information so the path computation engine can accurately calculate solution paths.

4. **Resource manager:** a database that stores all information about every reservation.

5. **Notification broker:** notifies the user or any other entity (e.g. a network management system) if the state of the reservation changes.

6. **Path computation engine:** determines if a solution path that meets all the constraints specified by the user exists within the network link topology graph.

7. **Device drivers (Path Setup System):** communicate directly with the network elements (e.g. routers, switches, etc) to set up and tear down the circuits as specified by the path computation engine for the duration and bandwidth requested by the user. This component understands vendor platform-specific technologies. This ability to abstract the underlying network element hardware is a critical aspect to making OSCARS vendor- and technology-agnostic and allows OSCARS to quickly assimilate new and emerging technologies such as OpenFlow.

8. **Workflow coordinator:** coordinates the workflow of the other components and user requests across network domains to facilitate the network reservation process.

Feedback

My group at UCSD is a member of the CMS experiment (Compact Muon Spectrometer) at the Large Hadron Collider in Geneva, Switzerland. As part of that, we operate one of seven data analysis centers,

referred to as T2 centers, for CMS in the US. Our center receives data in near-real time from the experiment. We routinely look at data within hours after it is processed. On occasion, this involves pulling tens of Terabytes of data to UCSD on very short notice. WAN connectivity allows us to reach a maximum of 15Gbps into the center.

OSCARS is the technology deployed within ESnet (and Internet2) to establish the data circuits that US-CMS uses. The UCSD/FNAL data circuit is based on OSCARS technology, as are all the other US-LHC T1/T2 circuits. The US-LHC T0/T1 circuits are also partially made up of OSCARS circuits. The UCSD/FNAL connection is especially important for our research, as FNAL is the archival storage location of the experiment within the US. Roughly 40% of all the data of CMS globally is archived there. The UCSD/FNAL data circuit enables the UCSD Tier-2 to move very large amounts of CMS data in a predictive manner and without impacting the general campus network traffic. As the UCSD T2 center is the primary analysis center for CMS members from UCSD, UCSB, and UCR, the OSCARS technology is thus vital for researchers across multiple campuses within the UC system.

– Frank Wuerthwein, Professor of Physics, University of California, San Diego

OSCARS gives us a framework to develop and experiment with novel and creative path computation algorithms. OSCARS provides us with the means to test them in a pseudo production environment to figure out real world effects of the theoretical algorithms. Once we have tried them out it is then a straightforward matter to apply them to actual production use. Essentially, OSCARS provides us a vehicle for technology transfer into an actual production environment—where theory meets the real world. We have found OSCARS to be extremely useful in developing and testing survivability and path protection algorithms.

– Biswanath Mukherjee, Professor of Computer Science, University of California, Davis

As the host laboratory for ATLAS data, Brookhaven National Laboratory is responsible for the receipt and processing of raw data from the ATLAS detector, before it is sent around the world. BNL depends on OSCARS circuits to provide the necessary connectivity to Tier 0, Tier 1 and Tier 2 facilities, in order to collect and transport the considerable and increasing amount of data that the ATLAS detector generates. OSCARS guarantees bandwidth and deterministic results during network failures for our services, enabling BNL to have essential bandwidth available upon demand to handle our data transfer and management needs.

OSCARS brings the benefits of schedulable, reliable bandwidth to a much broader scientific community in addition to leading science projects by providing campuses and regional networks with the necessary inter-domain circuit provisioning technology at low cost, and making high throughput services available to many scientists throughout the campus that need them, in addition to the thousands of physicists using the LHC Tier2 and Tier3 centers.

By using OSCARS to access large data sets and learn how to operate and manage global-scale networked systems, U.S. educational institutions remain on the leading edge of research necessary to attract the best students and faculty from around the world. OSCARS is especially important to smaller sites that may not have access to crucial network and/or storage resources otherwise.

– Michael Ernst Director, RHIC & ATLAS Computing Facility, U.S. ATLAS Facility Manager

I am writing in light of JGI's experience in July of 2010, when our team worked with DOE's Energy Science Network (ESnet) using OSCARS to address an immediate and urgent need to extend our compute resources over a wide area network. With ESnet as the intermediary, we virtualized the JGI data center, bringing it to NERSC, the National Energy Research Scientific Computing Center in Oakland, CA, 30 miles

away from our Walnut Creek facility. To accomplish this high-speed data transfer, we needed guaranteed bandwidth. We requested an OSCARS circuit from ESnet to provide us with private network connectivity between NERSC and JGI compute resources. Engineers from JGI facilities collaborated with ESnet and NERSC to set up a virtual circuit within minutes, which allowed the service to be initiated within hours. The circuit was robust enough to carry the data load as long as needed. From the setup of July 29, 2010, to date, over 1 petabyte of data has traversed the OSCARS circuits, which are still in production use today. We can now use the capability that OSCARS provides not only to reliably reroute and stream heavy loads of data over a wide area network in extraordinary circumstances, but to enhance our general operations as well. This is an important advantage for data-intensive research institutions like ours that periodically require additional resources, or commonly engage in long-distance collaborations.

– Victor Markowitz, Chief Informatics Officer and Associate Director, Joint Genome Institute

OSCARS is one of the principal enabling network and distributed system technologies that makes it possible for hundreds of high energy physics groups throughout the world involved in the Large Hadron Collider (LHC) program to process and analyze massive datasets as they explore the high energy frontier… The ability to exchange data at rates in excess of 10 Petabytes per month, in the form of many simultaneous 1–10 gigabit/sec flows, and also some larger aggregate flows, is crucial to the process of data exchange among the sites, and thus to the entire data analysis and physics discovery program…It is OSCARS' ability to dynamically allocate bandwidth and create virtual circuits across multiple network domains spanning intercontinental distances that is essential to the success of this entire program… OSCARS is also instrumental at Caltech, where we create dynamic circuits both for the LHC program and for the LIGO project.

– Harvey B. Newman, Professor of Physics, CalTech

Using OSCARS has enabled us to research different cross-domain issues because of the availability of the code and the motivated collaboration of the developer team at Berkeley.

– Prof. dr. ir.Cees T.A.M. de Laat, System and Network Engineering Research Group
Informatics Institute, Faculty of Science, University of Amsterdam

Project Team Members

Chin Guok
Evangelos Chaniotakis
Andy Lake
Mary Thompson
Eric Lomax

Submitter's name, title, and contact information

Wendy Tsabba
CRD Communications
Lawrence Berkeley National Laboratory
wtsabba@lbl.gov
510-486-7584