ESnet
Energy Sciences Network

January 2006 Volume 1, Issue 1

www.es.net

Berkeley Lab

news

ESnet Plays Key Role in International Effort of Building Trust on Grids

When the International Grid Trust Federation (IGTF) was established in October 2005 during the 15th Global Grid Forum (GGF) in Boston, much of the groundwork for the project had been laid by ESnet over the past three years. With the establishment of the ITGF, users of Grid computing worldwide took a step closer to accessing computers and information in 50 countries in Asia, the Americas and Europe. The IGTF brings together Grid organizations representing Asia, the Americas and Europe that are working towards allowing scientific researchers to identify themselves to any Grid resource in the world with just a single online identity.

IGTF's members issue electronic certificates that allow scientists to use the Grid. The Grids protected by IGTF certificates include over 40,000 computer processors and petabytes of storage — equivalent to over one million DVDs. Making sure the owners of Grids trust each other's security procedures is key to letting researchers access all these resources.

"By establishing IGTF, we are seeing the fruition of the first steps ESnet and the European Grid took back in February 2002, when a researcher at Fermilab used an authenticating certificate created by ESnet to successfully transfer files to Imperial College and Lancaster University in the UK," said ESnet's Tony Genovese.

"By establishing IGTF, we are seeing the fruition of the first steps ESnet and the European Grid took back in February 2002, when a researcher at Fermilab used an authenticating certificate created by ESnet to successfully transfer files to Imperial College and Lancaster University in the UK,"

continued on page 2

ESnet Provides Bandwidth Boost for SC|05 Bandwidth Challenge Entries

SEATTLE, Wash. — ESnet provided critical support for two entries in the Bandwidth Challenge at the SC|05 conference. The High Energy Physics entry, which included participation by the Stanford Linear Accelerator Center (SLAC) and the Fermi National Accelerator Laboratory (FNAL), was the winning project in the Bandwidth Challenge. The second ESnet-support effort was the National Energy Research Scientific Computing (NERSC) Center's Tri-Challenge entry, which encompassed the HPC Analytics and StorCloud challenges in addition to the Bandwidth Challenge.

The NERSC application utilized several high performance ESnet resources, including the new Bay Area Metropolitan Area Network (BAMAN) and the Science Data Network (SDN). Because of the capabilities of these infrastructures, the NERSC application did not require any special network engineering, even with anticipated data rates in excess of 5 Gb/s, and the traffic was routed over production infrastructure without any problems.

The High Energy Physics application

continued on page 2

INTRODUCING *ESnet NEWS*

This is the first issue of *ESnet News,* a new publication highlighting accomplishments and services of the U.S. Department of Energy's ESnet (Energy Sciences Network), managed by Lawrence Berkeley National Laboratory. Initially, *ESnet News* will be distributed every three months via email and may be freely distributed. For more information about ESnet news, write to info@es.net.

ESnet Increases 10 Gb/s Circuits by 880 Percent in 2005-06

ESnet will install at least 44 new 10 Gb/s circuits and connections in 2005-06. Half of these were completed and in production by November 2005, and the remaining connections should go into production by the end of spring 2006. By comparison, only five 10 Gb/s links were in place in January 2005.

These new connections are part of a new architecture that will provide 20-40 Gb/s metro area rings to dually connect the Department of Energy's Office of Science labs and will provide a second national core network — the Science Data Network (SDN) — for high-throughput science data and collaboration.

A significant achievement in 2005 was deploying the San Francisco Bay Area Metropolitan Area Network (BAMAN). This network consists of dual 10 Gigibit Ethernet (GE) rings that connect eight Bay Area research sites and provides 30-40 Gb/s connectivity for the Joint Genome Institute, the Stanford Linear Accelerator Center, the National Energy Research Scientific Computing Center, and

continued on page 4

ESnet Helps Build Trust on Grids *cont.*

said ESnet's Tony Genovese. "We did this as part of the pilot for the Particle Physics Data Grid. Once the British sites and Fermilab recognized and accepted each other's certificates, the data transfer went smoothly. With last week's decision, we now have a global trust architecture."

The IGTF brings Grid-oriented organizations around the globe much closer to realizing the promise of Grids. Grids aim to harness the power of geographically dispersed computing resources, experimental facilities and research centers. Grid developers' goal is to provide seamless access to all the resources available. However, at present there are many independently operated Grids, spread throughout the world, and users able to work on one can't necessarily gain access to the others.

Fundamental to user access is user authentication — making sure that only those users who have the proper credentials are granted access to the resources. While this can be a significant challenge within a Grid, achieving agreement on how to provide this level of authentication between Grids has been an even bigger challenge. That's where the IGTF takes center stage. With the establishing of the IGTF, the foundation is laid for building a trusted basis for identity management, and a further step is taken towards global interoperability for scientific grids.

The IGTF is a federation of certification authorities or Grid policy management authorities (Grid PMAs), and the major Grid infrastructure projects that together define the policies and standards for Grid identity management. Comprising the three regional Grid policy management bodies, the Asia Pacific Grid PMA (APGridPMA), the European Policy Management Authority for Grid Authentication in e-Science (EUGridPMA) and the Americas GridPMA (TAGPMA), the federation today has 61 members and covers 50 countries and regions.

For more information on the IGTF, go to <http://www.gridpma.org/>.

ESnet Provides Bandwidth Boost for SC|05 Bandwidth Entries *cont.*

(http://www-iepm.slac.stanford.edu/monitoring/bulk/sc2005/hiperf.html) attempted network utilization of unprecedented scale (in fact, there is no current production network that could have carried all the traffic), and so required significant custom configuration, donations of circuits, and collaboration with other networks. An example of the collaboration that supported the HEP entry was between ESnet and the research-driven Ultra Science Net (USN).

The HEP entry utilized two 10 Gigabit USN circuits, one of which is a long-term connection between ESnet and USN for use by SLAC and others. Two other, separate 10 Gigabit paths across ESnet were used by the HEP Bandwidth Challenge entry — one traversed the ESnet BAMAN beginning at SLAC and connected to USN for transport to the SC|05 show floor. The other began at SLAC, traversed the BAMAN, and was then routed over the ESnet Science Data Network to Seattle, where it connected to SC|05.

Another link requiring multiple collaborators was built to carry traffic from Fermilab to SC|05 for the HEP Bandwidth Challenge entry. This path traversed dark fiber donated by IWire from Starlight to the Qwest Communications POP in Chicago, a circuit from Chicago to Sunnyvale donated by Qwest Communications, a cross connect donated by Level(3) Communications in Sunnyvale, the ESnet SDN from Sunnyvale to San Diego, routers and cross connects donated by the San Diego Supercomputer Center, and finally the TeraGrid carried the traffic from SDSC to the SC|05 show floor in Seattle.

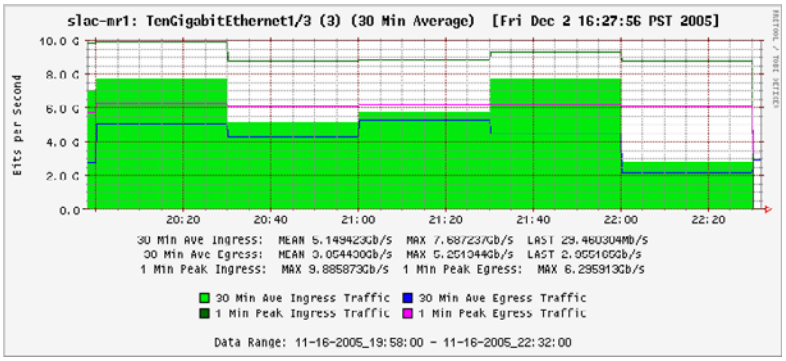

ESnet SDN traffic measured at SLAC

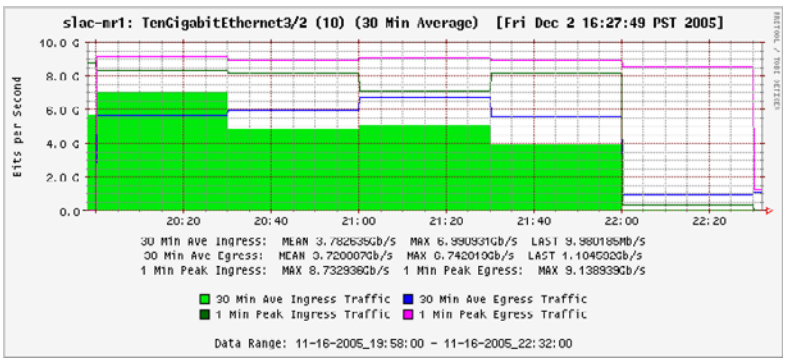

ESnet to USN traffic measured at SLAC

"These are examples of the collaborations necessary in the high performance computing and networking community" said ESnet Network Engineer Eli Dart, a member of the SC conference's SCinet team.

An example of the throughput achieved over ESnet can be seen in the utilization graphs of the two interfaces used by SLAC in the HEP Bandwidth Challenge entry. The first graph shows traffic measured at the ESnet SLAC router traversing the ESnet SDN path to SC|05. The second image shows traffic measured at the ESnet SLAC router traversing the Ultra-Science Net path to SC|05.

DOEGrids CA Enters Fourth Year of Operation — 3,500 Certificates Issued

DOEGrids Certificate Authority (CA) entered its fourth year of operation in December 2005 with about 3,500 active certificates, making DOEGrids one of the largest certification authorities in Grid computing. The growth of service and host certificates continues to outpace user certificates, as shown in the figure below.

DOEGrids is supporting development of new technology to improve the usability of DOEGrids PKI by its subscribers, administrators and relying parties. New scripts are under development in the community to automate aspects of the certification workflow, particularly for customers who need to certify hosts in a Grid computing cluster.

In collaboration with partners at LBNL and other research sites, ESnet is developing a credential store service for the Particle Physics Data Grid (PPDG). The credential store service, a kind of continent-wide MyProxy service, will provide network-based storage for Grid credentials and permit a more satisfactory and flexible roaming solution for PPDG users.

ESnet is also developing a short term certificate service with the capability of federating site authentication services and automatically issuing a short term Grid certificate based on a user's successful login to his site's authentication service. This certification service will offer a very easy roaming solution to a large subset of the DOE community, and open up Grid computing to a substantially larger population.

DOEGrids CA is valid until 2008, and these new techniques as well as others are under consideration for incorporation into the next generation of the DOEGrids CA, for which the planning process has just begun.

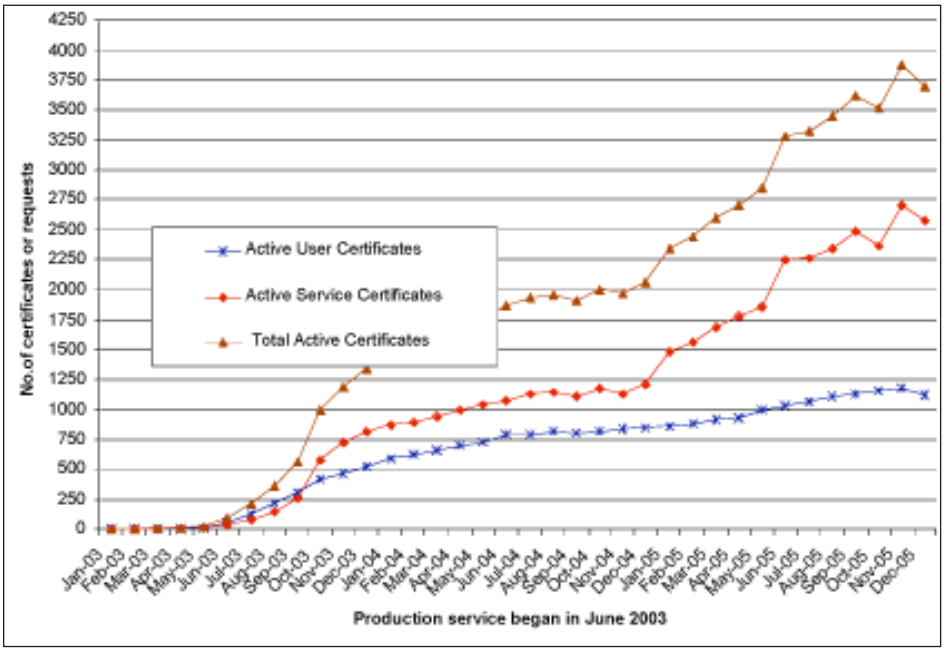

Growth of DOEGrids Certificate Authority

ESCC Update

By Scott Bradley, ESCC Chair

It has been a busy few months for the ESnet Site Coordinating Committee (ESCC). Our July '05 meeting was held as part of the ongoing Joint Techs conferences with Internet2 and Canada's advanced Internet development organization, CANARIE, in Vancouver, B.C. In addition to collaborating with our counterparts from the academic and scientific communities from both nations on general networking issues, the ESCC has been placing great focus on the unprecedented networking challenges being presented to us by the LHC community and Grid computing in general.

Never before has the DOE networking community had such a compressed implementation timeline between R&D efforts and production. This means that collaboration amongst ourselves and the community at large becomes even more crucial. To meet this need, the ESCC has created a series of technical focus groups, held both in person and virtually, to help ensure that as a community we are as responsive as possible to meet DOE's overarching networking requirements. Particular emphasis has been placed on the understanding of and participation in DOE-driven network research initiatives to ensure that our networking technology is as cutting-edge as possible.

Additionally, the ESCC website has been greatly enhanced (thank you ESnet!) to support rapid sharing of network-related files and utilities, and the ability to blog about them as well.

While the requirements for DOE networking continue to grow by leaps and bounds, the ESCC had more than kept pace to ensure we meet the challenges ahead.

ESnet Collaboration Services Update

ESnet Collaboration Services continues to expand the services that allow DOE Office of Science researchers, and their collaborators, the ability to meet and exchange information remotely as easily as if they were in the same location.

ESnet Collaboration currently has over 1,000 registered users worldwide of its voice, video, streaming, and data collaboration technologies which support such science initiatives as ATLAS, D0, CDF, ILC, CMS, ZEUS, Alcator C-Mod, OSG, DOSAR, LHC and others.

In 2005, streaming and H.239 services were added. For those who are unable to attend a meeting in person, they can use their browser and QuickTime or RealPlayer to watch the meeting via the stream. H.239 gives users the ability to share their computer desktop with anyone in the video meeting and can be used for presentations such as the one shown here from the second ESnet Collaboration Workshop (figure above). The 2005 ESnet Collaboration Workshop brought together customers and technical experts from the worldwide community to discuss needs and wants and hear presentations from vendors, researchers and users. The workshops culminate in a written document which helps guide ESnet Collaboration technology deployment to meet customer needs.

Remote workshop participants viewing slide presentation using H.239

The IP-video conferencing service (H.323) reached a milestone in November 2005 when it accounted for over 6000 port hours of use for the month. That translates to about 272 people-hours of meetings per workday for video only. If the 2000 hours of audio conferencing are added to that, ESnet supported a total of about 363 people-hours of meetings per day.

ESnet Collaboration staff continue to test, research and develop new technologies. One technology being tested is voice over IP (VoIP) as a possible addition to telephony-based audio conferencing. Because VoIP is less expensive than telephony, if the service is implemented it's expected to lead to greater use of audio conferencing.

10 Gb/s Connections Up 880 Percent *cont.*

Lawrence Berkeley, Lawrence Livermore and Sandia national laboratories.

ESnet also deployed the first three SDN circuits. These circuits provide 10 Gb/s service from San Diego to Sunnyvale, and Sunnyvale to Seattle, and an additional path between the IP hub and the SDN hub in Sunnyvale. These circuits replace 155 Mbps circuits, resulting in connections for laboratories in Southern California and Washington that are now 65 times faster.

In early 2006, ESnet will focus on deploying metro networks in Chicago, New York and Virginia-Maryland. The upgrades in Chicago will provide multiple 10 GE circuits to Argonne and Fermi labs, as well as connecting to Starlight and the ESnet backbone hub. This includes five 10 GE links to FNAL, which is a Tier 1 center for the Large Hadron Collider (LHC) project at CERN. The network in Long Island, NY, will provide two independent 10 Gb/s circuits to Brookhaven National Laboratory, which is also a LHC Tier 1 center. The network in the Washington, D.C. area will provide a single 10 Gb/s circuit to Jefferson Lab.

In addition to upgrading laboratory access links, ESnet is also improving network performance by upgrading many peering and cross-connections to 10 Gigabits to eliminate network bottlenecks between scientists at DOE laboratories, and instruments or collaborators at universities and research organizations world wide.

10 GB/S OR UPGRADED PEERING POINTS
- MANLAN
- MAX
- NGIX-EAST
- NGIX-WEST
- Starlight
- PNWGigapop
- PWAVE

NEW 10 GB/S BACKBONE LINKS
- New York to D.C.

NEW NETWORK INTERCONNECTS AT 10 GB/S
- USLHCnet
- UltraScienceNet

OTHER NEW 10 GB/S CONNECTIONS
- NASA Ames/ Columbia

Web-Based Services Support ESnet Users

ESnet provides a wide range of Web-based services to help users analyze and improve network performance. Here's a partial list of useful sites:

THE ESNET PERFORMANCE CENTER is a Web portal that enables users to run network tests between high speed UNIX-based hosts located at ESnet hub sites and to/from user end hosts. The focus of this tool is to assist in debugging network performance issues. See https://performance.es.net.

NETINFO is a Web site publishing ESnet network statistics. See http://netinfo.es.net.

PERFSONAR Traceroute Visualizer is a tool that leverages PerfSONAR Measurement\Archives to display link capacity and utilization at many points in a path that crosses multiple research and education networks. See https://performance.es.net/cgi-bin/level0/perfsonar-trace.cgi.

ESNET AVAILABILITY REPORTS provides detailed 12-month availability reports for each ESnet customer site. See http://calendar.es.net/cgi-bin/AvailReport.pl.

PUBLIC MAINTENANCE CALENDAR lists all currently planned and archived network maintenance and outage notifications from the ESnet NOC, including information on associated services. See http://calendar.es.net/cgi-bin/pmcalendar.pl.

ESnet Provides Key Support for SC|05 Conference

When the annual SC conference on high performance computing and networking convenes every November, the convention center where the event is held becomes the best-connected site on Earth — for one week. This dramatic achievement is due to contributions of staff effort and expertise by many of the world's leading networking organizations, including ESnet. And when SC|05 opened in Seattle in November, with 540 gigabits per second (Gb/s) wide area connectivity, support provided by ESnet helped make the conference the most successful to date.

In fact, ESnet has participated in the conference every year, helping plan, build and operate the conference network (known as SCinet), and borrowing the bandwidth conference that participants require for demonstrations.

At SC|05, ESnet provided 45 Gb/s (billion bits/second) of bandwidth for its customers, mostly to provide capacity for the Bandwidth Challenge, a spirited competition to show the ability to effectively use the large amounts of bandwidth made available by the major high-performance research and education networks, such as ESnet (see related story on p. 1).

In February, ESnet surveyed its sites to determine how much bandwidth was needed. Three sites responded, with Fermilab requesting a 10 Gb/s connection, NERSC asking for 5 Gb/s and Stanford Linear Accelerator Center requesting 20 Gb/s.

Provisioning so much temporary bandwidth is quite a challenge. Qwest, the primary bandwidth provider for Fermilab, initially agreed to donate the 10 Gb/s circuit, but then had to withdraw due to capacity problems in Seattle. But Qwest could get the circuit to Sunnyvale, from which point ESnet's new Science Data Network (SDN) connected to San Diego Supercomputer Center, which then utilized a 10 Gb/s TeraGrid circuit to connect to Seattle.

Unable to secure its requested bandwidth from National Lambda Rail, SLAC contacted Ultra Science Net (USN), a DOE network research project, which agreed to carry SLAC traffic if ESnet could get it to them in Sunnyvale. SDN already had the capacity and the connections to both USN and SLAC, so this was an obvious fit.

The second SLAC circuit would be routed over the main ESnet connection to SC|05 using the San Francisco Bay Area MAN to Sunnyvale and the SDN connection to SC|05 from there.

After ESnet got commitments for the requested 35 Gb/s of bandwidth, Lawrence Livermore National Laboratory made a late request for a 10 Gb/s connection to demonstrate Open InfiniBand. Pacific Wave, a partnership of CENIC and Pacific Northwest GigaPOP, offered ESnet a circuit from their Sunnyvale location to SC|05. To make the final connection to LLNL, Qwest and Ciena stepped forward and engineered and installed a 10 Gb/s link over the existing infrastructure between LLNL and Sunnyvale.

After almost two years of planning, the final staging, building and operation of SCinet took place over a fast-paced three weeks.

Hot staging began Oct. 25 in Seattle, with equipment being assembled, powered and tested; interoperability issues identified and fixed; and several high-speed network connections brought up.

On Monday, Nov. 7, the SCinet team began building out the network, installing 48 miles of fiber optic lines for booth connectivity, building several network centers on the show floor, and bringing up the rest of the high-speed links, including the 10 Gb/s link running Open InfiniBand to LLNL, with the help of LLNL, ESnet, NLR and University of Washington engineers. The network was ready when the show opened Monday, Nov. 14.

The highlight for the networking staff was the Bandwidth Challenge, a competition to make the best use of the "fat pipes" into the show. The high energy physics entry, including SLAC and Fermilab, won the top prize with the help of the bandwidth provided by ESnet.

At the end of the week, all of the hardware was disconnected and packed up for return to the vendors, and fiber was recovered from the hall for testing and re-use next year.

One key contributor who won't be returning in 2006 is Kevin Oberman, ESnet's longtime member of the SC networking committee. He is stepping down, with Eli Dart taking his place on the committee.

"Kevin will truly be missed," said Dart, who has worked with SCinet for more than five years. "His experience and his expertise have contributed mightily to the success and growth of SCinet over the years."

SC|05 was the most successful in the 18-year history of the conference, drawing more than 9,700 attendees, and the most successful SCinet — its 540 Gb/s connectivity was greater than that of many nations — with no major failures and very little downtime. This is due to the contributions of many individuals and companies. ESnet wants to thank all of the companies and individuals who helped make the ESnet success possible including Ciena, Ton Hutton (SDSC), David Richardson and the whole crew from Pacific Northwest GigaPOP, Qwest, Level(3), Pacific Wave, and a host of outstanding engineering folks from those companies as well as the SCinet team.

DOE Network Researchers Get a Consistent View with New perfSONAR Network Monitoring

In collaboration with GEANT2 and Internet2 networks, ESnet has released the first working prototype of a new network monitoring architecture, provisionally named perfSONAR. The prototype is now providing information on link capacity and link utilization across the circuits of the Abilene, ESnet, GEANT, UNINETT, and SWITCH backbone networks. Access is being provided initially to a restricted group, but with the intention of enabling more widespread access once the tool is stable and proven.

The prototype represents a real breakthrough for monitoring across complex networks like ESnet — this is the first time that data from disparate network locations has been made available in a consistent manner. Until now, monitoring data has been presented in different ways depending on where in the networks it was collected from, requiring network managers to use different tools for each network from which they needed to collect monitoring data.

If successful, the prototype service will establish a new seamless monitoring method that wipes out such differences between interconnected networks, enabling a true end-to-end view of network performance.

The positive impact of this new development is already being felt by the EGEE (Enabling Grids for E-SciencE) project, which is participating in the prototype phase. EGEE is receiving data through the perfSONAR interface for analysis by its own prototype diagnostic tool. This will allow the EGEE project to make decisions about network use based on the amount of available bandwidth between its participating sites, and to troubleshoot data transfer problems more easily.

The development of the prototype has been made possible by close cooperation with research teams in Europe and North America. Engineers at ESnet, Internet2, GEANT, Cesnet, DANTE, DFN, FCCN, GRNet, PSNC, Nordunet, Renater and RedIRIS have been part of the systems development, and some are also deploying the system on their own networks. This means that networks on both continents are presenting the same data in a consistent manner. This cooperation represents a step towards a seamless end-to-end measurement service across the major research networks in the U.S. and Europe.

OSCARS Prototype Provides On-Demand Circuit Provisioning to Researchers

ESnet is developing and deploying a prototype service that enables on-demand provisioning of guaranteed bandwidth, secure virtual circuits within the ESnet production network. The service, called OSCARS (ESnet On-Demand Secure Circuits and Advance Reservation System), is designed specifically to meet the science discipline-driven network requirement of dynamically provisioned quality of service (QoS) paths, as laid out in two DOE Office of Science workshops in the past two years.

The new service is being developed in close cooperation with Internet2/Abilene and the U.S. research and education network community, and with Dante/GEANT and the European research and education network community. The collaboration ensures that the new network services will interoperate among DOE labs and universities and other research institutions. Other collaborators include Brookhaven National Lab (Terapaths: A QoS-Enabled Collabor-ative Data Sharing Infrastructure for Peta-Scale Computing Research), General Atomics (network QoS for magnetic fusion) and Stanford Linear Accelerator Center (Internet end-to-end performance monitoring).

As a result of a Birds-of-a-Feather discussion at Internet2's Joint Tech meeting last year, ESnet and Internet2 have been co-developing the code base for the reservation system. Both ESnet and Internet2 have deployed a working prototype of the circuit-based service since early spring of 2005.

In an effort to promote inter-domain interoperability, ESnet, Internet and GEANT are in discussions on defining a common service definitions document. Work is also being done on workflow schemes to facilitate inter-domain reservation requests.