

The ESLEA Circuit Reservation Software

A.C. Davenhall^{*},^a P.E.L. Clarke,^a N. Pezzi^b and L. Liang^a

^a*National e-Science Centre, 15, South College Street, Edinburgh, EH8 9AA, UK*

^b*Network Group, Information Systems, University College London, 5, Gower Place, London, WC1E 6BS, UK*

E-mail: clive@nesc.ac.uk

We describe the Circuit Reservation Software (CRS) developed by the ESLEA Project. Switched-circuit networks can offer improved performance over conventional packet-switched networks for applications with some demanding types of network requirements. However, if switched-circuit networks are used to connect scarce and expensive resources, such as supercomputers, then advanced reservations or bookings of circuits are necessary if the resources and networks are all to be used efficiently. We describe and discuss the CRS, a demonstrator system for making such advanced bookings.

*Lighting the Blue Touchpaper for UK e-Science - Closing Conference of ESLEA Project
March 26-28, 2007
Edinburgh*

^{*}Speaker.

1. Introduction

We describe the Circuit Reservation Software (CRS) developed by the ESLEA Project for making advanced reservations of dedicated network connections. This software has a distributed and flexible architecture that makes it applicable to a wide variety of types of network. Deployment and use of the software will provide valuable practical experience of reserving and using dedicated network connections.

The ESLEA (Exploitation of Switched Light Paths for e-Science Applications) Project[1]¹ is a collaboration comprising groups working in computer science and a variety of scientific disciplines who are all using high-capacity dedicated network connections. Specifically, all the groups are using UKLight,² the UK research and development network for switched-circuit networks. The aim of the project is to investigate ways of using these connections and to share the resulting expertise both within the project and more widely. The project runs from February 2005 to July 2007.

The packet-switched Internet has been incredibly successful. However, there are now numerous scientific applications that require high network performance which is difficult to achieve over the conventional, production network. One alternative is to use special-purpose connections dedicated to the flows, and the purpose of ESLEA is to investigate this technique.

Such circuit-switched networks have been around for many years. They have received renewed interest in recent times because of developments in optical network hardware, notably DWDM (Dense Wavelength Division Multiplexing) in optical fibres and all-optical switches. There are, however, many ways to to implement (or simulate) direct connections. For example, emerging protocol standards such as GMPLS (Generalised Multi-Protocol Label Switching; see for example [2]) offer the possibility of extending dedicated flows to campus end-hosts without the provision of special hardware or manual configuration.

Currently circuits are usually configured ‘manually’ and typically persist for at least a few weeks and usually much longer. Many applications require connections of a shorter duration. It is only practical to provide such ephemeral circuits if they can be configured automatically. Further, it should be possible to ‘reserve’ or ‘book’ them in advance. This facility is necessary both to allow work to be planned effectively and to allow the network connection to be provided at the same time as the scarce resources that it will connect, typically super-computers. The CRS has been developed as a prototype to investigate making advance reservations in circuit-switched networks. A previous paper presented our plans early in the Project[3]. The present paper describes and discusses the CRS.

2. The CRS Architecture

The purpose of the CRS is to create a dedicated connection between two end-hosts, one of which is local and the other geographically remote. The CRS adopted the architecture of the EGEE (Enabling Grids for E-Science) BAR (Bandwidth Allocation and Reservation) Project³ which had

¹See URL: <http://www.eslea.uklight.ac.uk/>

²See URL: <http://www.uklight.ac.uk/>

³See URL: <http://egee-jra4.web.cern.ch/EGEE-JRA4/>

developed similar software for making reservations of guaranteed bandwidth between two end-hosts using QoS (Quality of Service) mechanisms operating in conventional networks. This architecture is general, robust, flexible and extensible.

The route of the circuit between the two end-hosts is considered to comprise three segments: the local campus network to which the local end-host is connected; the WAN (Wide Area Network) connecting the local campus network to the campus network where the remote host is located and the remote campus network to which the remote host is connected. The EGEE BAR architecture (see Figure 1) is distributed and has separate components for controlling each of these segments. In outline the architecture is as follows. Requests to create a new reservation are initiated by some client software, which can be either an interactive ‘reservations browser’ driven by a user or some piece of Grid middleware (usually called ‘higher layer middleware’ or HLM) acting autonomously. To make a reservation the client invokes the top-level component of the architecture, which is also called a BAR. The BAR component does not itself configure any network hardware. Rather, it controls subsidiary components which configure the hardware. It also maintains a database of existing reservations.

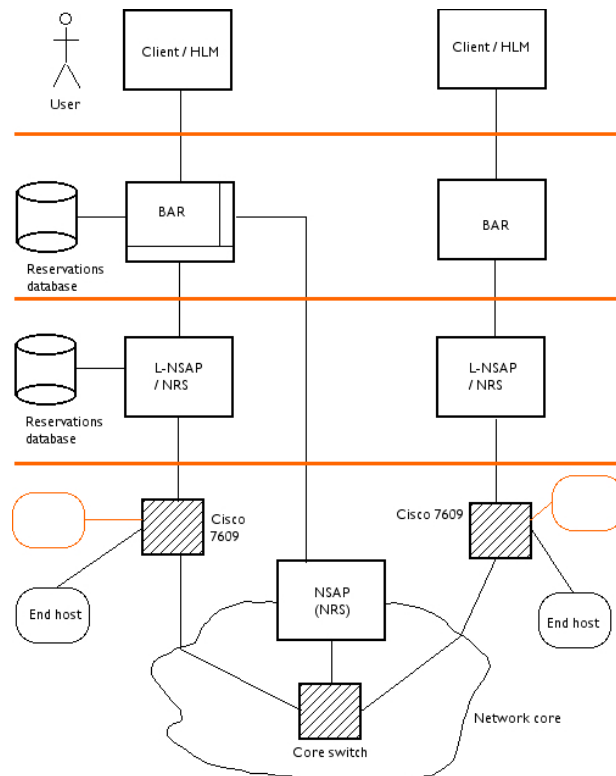


Figure 1: The CRS architecture

When a request to make a new reservation is received the BAR asks the individual subcomponents whether they can handle the request. If all can do so the reservation is accepted and added to the database. If the reservation cannot be accepted the user is informed. The individual components

are the L-NSAP and the NSAP. The L-NSAP (Local Network Services Access Point) configures the local campus network, and the NSAP (Network Services Access Point) configures the WAN.

The BAR directly invokes the L-NSAP that configures the local campus network, but it does not directly invoke the L-NSAP that configures the remote end-site. Rather, it communicates with the BAR responsible for the remote end-site, which in turn invokes its own L-NSAP.

3. The CRS Implementation

Each of the CRS components (see Figure 1) is implemented as a Web service written in Java and running on a Linux PC. The interfaces to the Web services and some of the code were adopted from the EGEE BAR Project. The L-NSAP and NSAP components were also not written *ab ovo*, but rather, are modified versions of the NRS (Network Reservation System)[4]⁴ developed by Saleem Bhatti *et al.* The CRS only uses items of external software that are usually available free of charge, such as Apache Tomcat,⁵ Apache Axis⁶ and the databases HSQLDB⁷ and PostgreSQL.⁸

Though the CRS architecture is flexible and general, the current *implementation* is quite limited in some respects. Switched-circuit connections to the participants in the ESLEA Project were provided across the UKLight network, which does not support the automatic configuration of circuits. We connected a Cisco 7609 switch/router to the UKLight access points at suitable institutions (these machines were available from the earlier MB-NG Project⁹). The 7609s simulate the LAN connecting the end-host to the WAN. Switching inside UKLight was simulated using an additional 7609 connected so as to be topologically inside UKLight whilst remaining physically outside it. The edge 7609s were controlled by the L-NSAP components and the simulated UKLight switching by the NSAP component.

4. Discussion

The major development of the CRS is complete. It has been tested on a small test network of UKLight circuits. Trial deployments with two collaborating groups in the ESLEA Project are continuing. However, we have already demonstrated the practicality of the basic idea of configuring switched-circuit networks from an advance booking system.

Further work that we anticipate includes collaborating with the UCLP¹⁰ (User Controlled Lightpaths) group to produce an NSAP that interfaces to the UCLP software for creating switched lightpaths across an optical network. We are also collaborating with the developers of the HARC (Highly-Available Robust Co-scheduler)[5, 6]¹¹ co-allocation software to allow our network reservations to be automatically co-scheduled with reservations of other resources.

⁴See URL: <http://www.cs.ucl.ac.uk/staff/S.Bhatti/grs/>

⁵See URL: <http://tomcat.apache.org>

⁶See URL: <http://ws.apache.org/axis/>

⁷See URL: <http://hsqldb.org/>

⁸See URL: <http://www.postgresql.org/>

⁹See URL: <http://www.eslea.uklight.ac.uk/mb-ng>

¹⁰See URL: <http://www.uclp.ca/>

¹¹See URL: <http://www.cct.lsu.edu/~maclaren/HARC/>

High capacity dedicated connections are likely to remain a scarce resource for at least the near future, as are the super-computers and similar devices which typically act as the sources and sinks for the data that flow through them. It seems likely that there will be a continuing need to reserve and schedule these scarce and expensive resources in advance. We have demonstrated the basic viability of making advanced reservations over a circuit-switched network.

Acknowledgments

We are grateful to Saleem Bhatti and Richard Smith for discussions about NRS; to Charaka Palansuriya, Kostas Kavoussanakis, Florian Scharinger and Alistair Phipps for discussions about the EGEE BAR effort and to Stefan Zasada, Jon MacLaren and Stephen Pickles for discussions about HARC and co-scheduling.

For assistance in setting up test circuits we are grateful to UKERNA, particularly John Graham and David Tinkler, and to Sam Wilson and his colleagues of EUCS. Stephen Kershaw and Barney Garrett assisted with setting up end-hosts.

ESLEA was funded by the UK Engineering and Physical Science Research Council, with additional contributions from the Medical Research Council and the Particle Physics and Astronomy Research Council.

References

- [1] C. Greenwood, V. Bartsch, P. Clarke, P. Coveney, C. Davenhall, B. Davies, M. Dunmore, B. Garrett, M. Handley, M. Harvey, R. Hughes-Jones, R. Jones, M. Lancaster, L. Momtahan, N. Pezzi, S. Pickles, R. Pinning, A. Simpson, R. Spencer, R. Tasker, *Exploitation of Switched Light Paths for e-Science Applications (ESLEA)*, in S.J. Cox, D.W. Walker (eds), *Proceedings of the UK e-Science All Hands Conference 2005*, Engineering and Physical Sciences Research Council (2005), CD-ROM.
- [2] A. Farrel, I. Bryskin, *GMPLS: Architecture and Applications*, Morgan Kaufmann, Amsterdam (2006).
- [3] A.C. Davenhall, P.E.L. Clarke, N. Pezzi, *The ESLEA Control Plane Software*, in S.J. Cox, D.W. Walker (eds), *Proceedings of the UK e-Science All Hands Conference 2005*, Engineering and Physical Sciences Research Council (2005), CD-ROM.
- [4] M. Rio, A. di Donato, F. Saka, N. Pezzi, R. Smith, S. Bhatti, P. Clarke, *Quality of Service Networking for High Performance Grid Applications*, *Journal of Grid Computing*, **1** (2003) 329-343.
- [5] J. MacLaren, *Co-allocation of Compute and Network resources using HARC*, in proceedings of *Lighting the Blue Touchpaper for UK e-Science: closing conference of ESLEA Project*, PoS(ESLEA)015 (2007); these proceedings.
- [6] J. MacLaren, M. McKeown, S. Pickles, *Co-Allocation, Fault Tolerance and Grid Computing*, in S.J. Cox (ed), *Proceedings of the UK e-Science All Hands Conference 2006*, Engineering and Physical Sciences Research Council (2006), CD-ROM.