

Exploitation of Switched Lightpaths for e-Science Applications (ESLEA)

C. Greenwood^g, V. Bartsch^{bd}, P. Clarke^g, P. Coveney^b, C. Davenhall^g, B. Davies^e, M. Dunmore^e, B. Garrett^c, M. Handley^b, M. Harvey^b, R. Hughes-Jones^f, R. Jones^e, M. Lancaster^b, L. Momtahan^h, N. Pezzi^b, S. Pickles^f, R. Pinning^f, A. Simpson^h, R. Spencer^f and R. Tasker^a

^aCCLRC Daresbury Laboratory, UK; ^bUniversity College London; ^cUniversity of Edinburgh; ^dFermilab National Accelerator Laboratory, USA; ^eLancaster University; ^fUniversity of Manchester; ^gNational e-Science Centre, Edinburgh; ^hUniversity of Oxford.

Abstract

A major imperative facing the research community is the sharing of high performance computing resources to meet data transfer demands of the distributed research paradigm. In recognition of the needs of high-end computer applications, UKLight, the UK's first national switched circuit optical network, was established in 2004. UKLight provides participating research centres throughout the UK with the infrastructure necessary to develop optical networking and the applications that will use such networks. It also facilitates collaboration on global research projects by providing access to the fast growing international optical network.

ESLEA is the first widely scoped project to exploit UKLight for a range of scientific applications, including Particle Physics, Radio Astronomy, Computational Science and e-Health. Recent work within these applications indicates that UKLight can reduce data transfer times from days to hours compared to packet routed networks. An overview of ESLEA and the pilot applications is provided and progress to date is discussed.

1. Introduction

The ESLEA project [1] intends to demonstrate that the needs of data-intensive applications are best met by optical networks ("lightpaths") with switched point-to-point circuits, rather than by the traditional and costly routing of information at the packet level. Circuit switching is a process of connecting two points in a network along a dedicated end-to-end path which remains in effect until the communication is terminated. Once the link is broken, a new circuit can be created. The UKLight circuit-switched network provides ESLEA with a suitable infrastructure.

UKLight [2] is an optical R&D network, managed through UKERNA [3] as a complement to the SuperJanet4 production network. It comprises a 10Gbit/s backbone to selected points in the UK and connects to global optical networks via 10Gbit/s links to Chicago (via StarLight) and Amsterdam (NetherLight) - see Figure 1. Multiplexing technology, based on time slots, supports multiple logical circuits so that projects can run concurrently; generally, 1Gbit/s streams are being used by ESLEA.

This paper provides an overview of the ESLEA project, which includes both capability

development and exploitation tasks. The capability component is common across the project and represents tasks that would otherwise be replicated if pilot applications were launched independently. The exploitation component involves integrating the experimental and analytical framework of pilot applications into UKLight to demonstrate the feasibility and benefits of high-bandwidth switched lightpath circuits. Project progress to date is reviewed for all tasks.

2. Capability development tasks

These tasks involve provision of support for the technical preparations needed to integrate each pilot application with UKLight as well as ongoing network engineering services. In addition, development of control plane software and network protocol research are being undertaken to help pilot applications exploit UKLight effectively.

2.1 Control plane software (CPS)

The dynamic provisioning of switched lightpaths in response to application requests necessitates the development of new control methodologies to create, remove and control end-to-end circuits across domain boundaries.

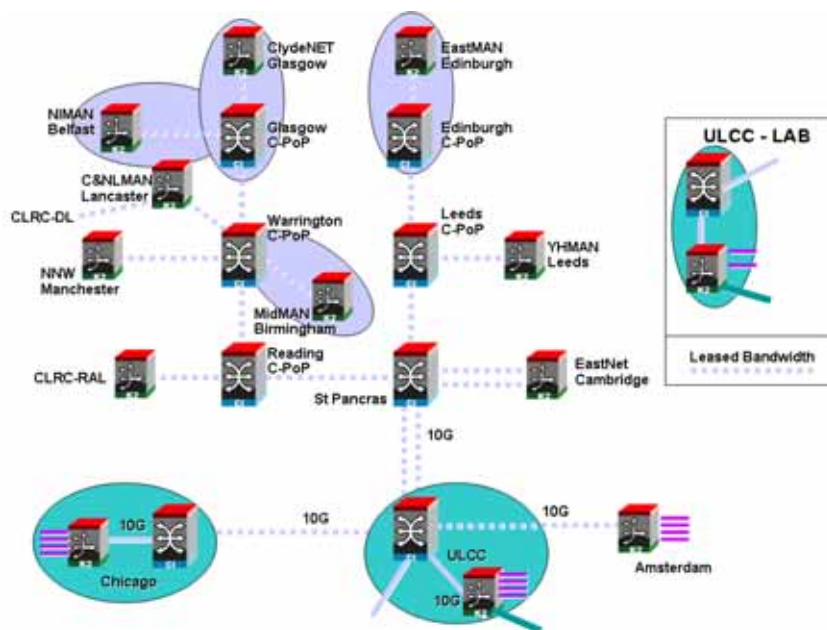


Figure 1. UKLight Infrastructure (courtesy UKERNA)

Initially, the emulation of switched lightpaths will be performed using statically configured Cisco 7609 switch/routers. It will not be possible for ESLEA engineers to perform any switching in the UKLight core, but it is planned that an optical switch outside the core will emulate such switching later in the project.

CPS is being developed so that end-users from the pilot applications will be able to reserve circuits in advance.

2.2 Transport protocols

Initial exploitation of UKLight infrastructure will use well-established transport mechanisms such as TCP. However, research will be conducted into the use of High Speed TCP and DCCP for data transport in hybrid optical networks, i.e. where access over conventional IP networks is combined with guaranteed, high bandwidth dedicated channels. The pilot application teams will be assisted to deploy these implementations and the experience gained will be used to refine and modify the algorithms and/or implementations.

3. Exploitation tasks

A number of e-Science applications have been identified with requirements that might be difficult, or impossible, to deliver with today's IP-based research networks. These applications, described below, have been chosen to pilot the use of UKLight within the ESLEA project.

3.1 CDF (Collider Detector at Fermilab)

This particle physics experiment takes data from proton - anti-proton collisions at Fermilab, Chicago. Currently, most data analysis takes place at Fermilab owing to the time overhead in copying data to remote sites. However, demand for resources at Fermilab is increasing and so rapid distribution of data for remote analysis and re-processing is required. The recent commissioning of optical networks in the USA (StarLight), UK (UKLight) and Netherlands (NetherLight) makes it feasible for Europe to host a significant fraction of the CDF data and to provide CPU resources for its analysis worldwide.

The CDF application within ESLEA aims to demonstrate that guaranteed high-bandwidth transfers of data between UK and Fermilab, via the UKLight and Starlight networks, can enable the UK to play a significant role in this analysis and reprocessing work.

3.2 ATLAS (A Toroidal LHC Apparatus)

ATLAS is one of four detectors being built at the Large Hadron Collider (LHC), CERN. Protons will be accelerated and steered to collide in the middle of the ATLAS detector. The debris of the collisions is expected to reveal fundamental particle processes.

As part of preparations for the start of physics collisions in 2007, ATLAS scientists are planning to generate and transfer increasing amounts of data in a series of Service Challenges. The aim of the Service Challenges

is to test the infrastructure needed for full LHC data taking and analysis.

The ATLAS application within ESLEA aims to demonstrate high-bandwidth transfers of data between Rutherford Appleton Laboratory (RAL) and CERN, via UKLight and NetherLight, as well as from Lancaster to RAL and Lancaster to Manchester using UKLight.

3.3 Very Long Baseline Interferometry (VLBI)

VLBI is used by radio astronomers to obtain the high resolution images of cosmic radio sources. VLBI experiments record data at separate sites with high-precisions timestamps and then each site ships tapes or disks holding this data to a central site where correlation is performed. The most advanced recording system in use today, the Mark V disk, has a data rate of 1 Gbit/sec.

High speed transfers of telescope data are planned between observatories at Jodrell Bank and Dwingeloo, via 1 Gbit/sec UKLight and SURFnet connections, for correlation in real time [or in quasi-real time after local buffering]. The aim is to show that switched lightpaths will provide the bandwidth needed to exploit fully the latest VLBI equipment, such as the Mark V disk. In addition, high-bandwidth transfer of data between the UK and Haystack correlator at MIT, via UKLight and StarLight, is planned.

3.4 RealityGrid

RealityGrid applications represent areas of High Performance Computing (HPC) with bandwidth intensive requirements for computational steering, visualisation and high performance data transfer.

The scientific domain of RealityGrid is the realistic modeling and simulation of complex condensed matter at the meso- and nano-scale levels as well as the discovery of new materials.

The HPC application within ESLEA aims to demonstrate high-bandwidth transfers of data across a Grid combining UK and TeraGrid HPC resources, linked by UKLight and StarLight.

3.5 Integrative Biology (IB)

The pilot IB application will demonstrate the potential of circuit switched optical networks to deliver digital, high resolution medical images. It involves the development of computer models of heart function, ranging in scale from molecular level to whole organ. This research opens up the prospect of using computer models to develop new treatment regimes and drugs to counter heart disease.

Bandwidth-intensive activities such as computational steering, performance control,

collaborative visualisation, data curation and data mining will be undertaken using UKLight connections between HPC facilities at Daresbury Laboratory, Manchester and RAL.

4. Discussion

The project is in its early stages and some applications have yet to be connected to UKLight. The following links show the ESLEA project's requirements for UKLight connectivity in 2005 (those connections yet to be made are shown in italics):

- 1*1.0 Gbit/s UCL - Fermilab (CDF)
- 2*1.0 Gbit/s RAL - CERN (ATLAS)
- 1*0.5 Gbit/s RAL - Lancaster (ATLAS)
- 1*1.0 Gbit/s Manchester - Dwingeloo (VLBI)
- 1*1.0 Gbit/s Manchester - Haystack (VLBI)*
- 1*0.5 Gbit/s UCL - Lancaster* (ReG)*
- 1*0.5 Gbit/s UCL - Manchester (ReG)*
- 1*0.5 Gbit/s UCL - TeraGrid (ReG)*
- 1*0.5 Gbit/s Manchester - Lancaster* (ReG)*
- 1*0.5 Gbit/s Manchester - TeraGrid (ReG)*
- 1*0.5 Gbit/s Lancaster* - TeraGrid (ReG)*
- 1* 1.0 Gbit/s RAL - Lancaster* (IB)*

* For onward connection to Daresbury Laboratory

The research into transport protocols is scheduled to commence during summer 2005, but work has already begun on developing control plane software. A simple GUI-based Java program has been implemented to configure connections in a Cisco 7609 switch/router in real-time [4]. In addition, Network Resource Scheduling software [5] has been converted to create switched-circuit connections. Thus, it is now possible to create switched circuits between ESLEA end-hosts placed around the edge of UKLight.

All exploitation applications have commenced and their progress to date is summarised below:

- **CDF.** A dedicated connection from UCL to Fermilab, via UKLight and StarLight, was established in late 2004. This link is currently undergoing integrity and bandwidth tests with the aim of achieving 1 Gbit/s line speed over several days [6]. Tests using the iperf network bandwidth monitoring tool show that the link can already deliver a data rate of approx. 800 Mbit/s.
- **ATLAS.** Two 1 Gbit/s links from RAL to CERN were used in Service Challenge (SC) 2 during March-April 2005. Over a 10 day period, an aggregate data transfer rate to disk of 576 Mbit/s was achieved, with a total volume of approximately

68 TBytes of disks space filled. SC3 is scheduled to finish at the end of July 2005. The UKLight link from Lancaster to RAL was connected in June 2005, with initial maximum transfer rates of 290Mbits/s being expected. The link is undergoing tests to commence transfer of files and both experimental and site software are being configured for SC3 [7].

- **VLBI.** Tests in the second quarter of 2005 using UDPmon on the Manchester - London (MB-NG test bed) - Amsterdam (UKLight) - Dwingeloo (SURFnet) link achieved a line rate throughput between Manchester and Dwingeloo (up to 999Mbit/s) for UDP packets of more than 1000 bytes. No packet reordering was seen but high packet loss occurred for small to medium size packets, probably in the end host [8].

In late May 2005, Jodrell Bank was connected to the MB-NG router in Manchester; rates of approximately 600 Mbit/s were obtained.

UKLight connected to the Warrington node during June 2005. UDPmon tests show that line rates (999 Mbit/s) can be obtained with no reordering and no packet loss for large packets of more than 1000 byte [9].

- **RealityGrid.** At the EPSRC e-Science meeting, Edinburgh, in April 2005, a 1 Gbit/s connection between Manchester and TeraGrid enabled real time visualisation and steering of a running lattice-Boltzmann simulation (of binary fluid flow in porous media) to be demonstrated successfully. StarLight and UKLight networks facilitated the transfer of approximately 0.6 Tbytes of data in less than a night, with up to 30 streams opened concurrently. Although the maximum bandwidth available was limited by a 100Mbit/s network card connecting two machines at UCL, the same transfer would have taken days on a normal network [10].
- **e-Health.** The application awaits connection to the UKLight link between Daresbury and RAL laboratories. Once connected, it is anticipated that simulation model runs will require TeraBytes of simulation data to be transferred between the computation site at Daresbury and the storage facility at RAL [11]. It is planned to use this link for visualisation of simulations and development work has already started on the software that will be required.

5. Summary

The ESLEA project intends to demonstrate the benefits of switched lightpaths by running data-intensive pilot applications over UKLight.

The RealityGrid application has already shown that UKLight can reduce data transfer times significantly compared to traditional packet routed networks. It is intended that ATLAS, VLBI, RealityGrid and Integrative Biology applications will provide further demonstrations of the benefits of switched lightpaths at forthcoming events in 2005. The CDF application is scheduled to commence transfer of simulated datasets from UCL to Fermilab by the end of 2005.

So far, much useful experience has already been gained in linking applications to UKLight and early indications are that this high bandwidth facility will significantly enhance the UK's e-Science research capability.

Acknowledgements

The authors thank David Salmon and John Graham for their valuable contribution to running the pilot applications over UKLight.

ESLEA is funded by EPSRC, MRC and PPARC.

References

- [1] <http://www.eslea.uklight.ac.uk>
- [2] <http://www.uklight.ac.uk/>
- [3] <http://www.ukerna.ac.uk/>
- [4] A.C. Davenhall, P Clarke and N Pezzi. The ESLEA Control Plane Software, *Proceedings of the 4th All Hands Meeting, Nottingham*, Sept 2005.
- [5] <http://www.cs.ucl.ac.uk/staff/S.Bhatti/grs/>
- [6] V Bartsch, P Clarke, B Garrett, C Greenwood, M Lancaster and N Pezzi. CDF's Utilisation of Switched Circuit Optical Networks. *Proceedings of the 4th All Hands Meeting, Nottingham*, Sept 2005.
- [7] B Davies, Lancaster University, *Pers. Comms*, June 2005.
- [8] R Spencer and R Hughes-Jones, University of Manchester, *Pers. Comms*, June 2005.
- [9] R Spencer, R Hughes-Jones and S Casey, High Rate Internet Data Transfer for eVLBI, *Proceedings of the 4th All Hands Meeting, Nottingham*, Sept 2005.
- [10] M Harvey, University College London, *Pers. Comms*, June 2005.
- [11] L Momtahan and A Simpson, Exploitation of Switched Lightpaths for e-Science Applications: an e-Health mini-project, *Proceedings of the 4th All Hands Meeting, Nottingham*, Sept 2005.