eScience Applications on the SURFnet RE Network

Jeroen van der Ham, Paola Grosso and Cees de Laat
System and Network Engineering Research Group, University of Amsterdam
{*vdham—p.grosso—delaat*}*@uva.nl*

Abstract: The hybrid network paradigm pioneered by SURFnet provides the capacity demanded by modern eScience applications. This contribution presents ongoing research and developments to integrate photonic networks in the application middleware stack.

1. Introduction

The architectures of national research and education networks (NRENs) are in continuous evolution. The primary goal is to optimally support the communication requirements of the connected students and researchers. For this reason NRENs have always been at the forefront in terms of adoption of new technologies as the user demands have often been cutting edge, both in terms of total required bandwidth and quality of service of the network, as well as in terms of global services.

The SURFnet network, with its advanced backbone and its hybrid network model, is since years one of the most advanced NRENs worldwide. An important driver for innovation in SURFnet is the support of eScience applications; in particular the definition of network models suitable for path finding, the research on QoS-aware workflows, and the creation of services for inter-domain dedicated circuits are key components in the realization of this infrastructure. In the following sections we provide brief summaries of these issues and their relation to the SURFnet network.

2. SURFnet Network

SURFnet is the NREN of the Netherlands. It has a fiber backbone across the whole country, serving all universities and colleges, see figure 1. This backbone is lighted using DWDM technology, which provides the basis for the creation of a hybrid network; one wavelength is used for the internet connectivity, the others are available for direct circuits between endpoints either in a static form known as Optical Private Networks, or as dynamic lightpaths.

SURFnet is uniquely positioned in the international network backbone. It was the first Internet connection point in Europe, and throughout the years, Amsterdam has maintained its position as an important exchange point for Internet traffic. SURFnet has pioneered the concept of open lightpath exchanges with the creation of Netherlight [1]. Netherlight is an important hub with over a dozen trans-atlantic connections, connections to the European GEANT network, and direct peerings with many other NRENs in Europe. Prime examples are the new dark fiber connections to CERN and Nordunet.

SURFnet and the University of Amsterdam are among the founders of the Global Lambda Integrated Facility (GLIF) [2], a virtual organization promoting the paradigm of lambda networking where open lightpaths exchanges such as Netherlight are coupled together.

3. eScience on SURFnet

The hybrid network model, and the excellent global connectivity puts the SURFnet network in a unique position for enabling eScience applications. eScience is, simply said, science that makes use of modern ICT infrastructures to perform computation and experiments. Support for cooperation and decentralization play an essential role: scientists cooperate among institutions to maximize the sharing of expertize and human resource, while data is distributed across domains, to leverage computing, storage or instrumentation resources. eScience infrastructures provide the tools, the services and the support to allow researchers to work according to this new model.

A first attempt for the integration of dynamic photonic network connections and applications was investigated in the StarPlane project [3]. The Distributed ASCI Supercomputer (DAS) consists of a compute cluster distributed over

Fig. 1. The dark fiber infrastructure of SURFnet

four locations in the Netherlands. In the StarPlane project one WDM band was made available to dynamically create different network topologies between these locations. The coupling of Grid middleware to the network reservation and provisioning system allowed for the first time a close synchronization of the application needs to the underlying photonic paths.

Currently many more eScience applications exploit the advantages of the SURFnet hybrid architecture; the CineGrid and SCARIe project well illustrate this synergy between advanced network services and application needs. For the CineGrid project [4] SURFnet transports high-quality digital media material, of resolution of 4K and higher, to various locations in the country and internationally. The data are located in large storage nodes, with capacity of hundreds of TBs, and are transported to visualization facilities for viewing and analysis. In this case the major requirement is the availability of high-bandwidth paths without competing traffic that would vary the jitter and latency of the packets arrival, which in turn would result in a deteriorated video stream.

For the SCARIe project [5] data from radio telescopes is sent to correlation centres around the country, to be analyzed in very long baseline interferometry (VLBI) experiments. Here the possibility of linking many remote sites with ad-hoc connections allows faster computation, which results in better utilization of the expensive telescope time for radio astronomers.

There are two main functionalities that eScience applications require to optimally use the advanced services provided by the hybrid network model: schedule lightpaths and compose network-aware workflows.

Scheduling of optical circuits requires a priori knowledge of the network topology. The Network Description Language (NDL) [6, 7] defines an ontology to describe (optical) networks. This allows the engineers to have a detailed and accurate view of the network, and use it to find paths. The strength of this approach is that it provides a richer semantic for the determination of feasible paths. The topology description can not only span multiple network layers, but also allows the integration of network information with semantic descriptions of other resources in the whole eInfrastructure.

Scientific workflows rarely include the network QoS in their composition and execution cycle, but this is an essential component to fully utilize the enormous network potential of hybrid architectures. The University of Amsterdam (UvA) has researched the inclusion of network resources and the associated QoS metrics in these processes. The solution proposed is an RDF based schema for describing the data and QoS requirements in an abstract workflow, which can be consumed by proper agents during the various moment of the scientific workflow lifecycle [8].

4. Inter-Domain Networking

Dynamic provisioning of lightpaths will allow for real dynamic eScience in the future. Even more powerful will be creation of dynamic paths across domains.

The network of lightpaths used by the LHC experiments at CERN shows what future eScience can bring us. The data coming out of the LHC experiments are being distributed worldwide over the optical networks, first to Tier-1 institutes, then to Tier-2, et cetera. These data streams are in the order of Petabytes per month. The current LHC distribution network comprises of static between the institutes around the globe. However, with the rise of dynamic optical circuit networking these inter-domain circuits can become dynamic.

The LHC experiments have taken years to construct, and will continue for the next several years. The new dynamic lightpaths will make it possible also for short lived experiments to use dynamic lightpaths. An example was already shown where two supercomputers were hooked up together to work on a simultaneous simulation of the universe. The lightpaths there took several days to be provisioned and configured correctly. Making it possible to do this automatically and dynamically will allow for many new interesting and new research opportunities [9].

Still, inter-domain dynamic lightpaths are in their infancy: lightpaths are either completely configured by hand or are manually inserted into the provisioning software of each domain. The Network Service Interface working group (NSI-wg) of the OGF is working on a standard interface for inter-domain circuit provisioning. Implementations of this experimental interface is already ongoing in the openDRAC software used in SURFnet and demonstrations of real dynamic provisioning have been given at SuperComputing 2010.

For inter-domain dynamic lightpaths it is necessary to share connectivity information with connected networks. NDL allows domains to provide a clear description of their entire network, which can be hyperlinked to describe connections to neighboring networks. Sharing the complete topology may not be desirable, because of security, business, or scalability reasons. It is also possible to share an aggregated view of the network. These aggregations can range from a full mesh with border nodes with some details of internal connectivity to representing the whole domain as a single node. However, this latter approach can have a dramatic effect on the success rate of inter-domain pathfinding [10], leading to a large number of false-positives. The full-mesh aggregation approach maintains the most relevant information for inter-domain pathfinding. The performance loss with this aggregation method for inter-domain pathfinding is almost negligible.

5. Future work and conclusions

Many emerging themes are driving the current developments in SURFnet and are the focus of future work. For example, the use of green technologies or the emergence of virtual resources are interesting upcoming research areas which will push further the integration of hybrid network in the intelligent composition and management of eInfrastructures.

References

1. NetherLight, "http://www.netherlight.net/," .
2. Global Lambda Integrated Facility (GLIF), "http://www.glif.is/," .
3. P. Grosso, D. Marchal, J. Maassen, E. Bernier, L. Xu, and C. de Laat, "Dynamic photonic lightpaths in the starplane network," Future Generation Computer Systems **25**, 132–136 (2009).
4. CineGrid, "http://www.cinegrid.org/," .
5. N. Kruithof and D. Marchal, "Real-time software correlation," in "INGRID Workshop," (2008).
6. J. van der Ham, F. Dijkstra, F. Travostino, H. Andree, and C. de Laat, "Using rdf to describe networks," Future Generation Computer Systems, Feature topic iGrid 2005 (2006).
7. J. van der Ham, P. Grosso, R. van der Pol, A. Toonk, and C. de Laat, "Using the network description language in optical networks," in "Tenth IFIP/IEEE Symposium on Integrated Network Management," (2007).
8. Z. Zhao, P. Grosso, R. Koning, J. van der Ham, and C. de Laat, "An architecture including network qos in scientific workflows," in "International Conference on Grid and Cloud Computing (GCC 2010)," (2010).
9. S. P. Zwart, T. Ishiyama, D. Groen, J. M. Keigo Nitadori, C. de Laat, S. McMillan, K. Hiraki, S. Harfst, and P. Grosso, "Simulating the universe on an intercontinental grid of supercomputers," IEEE Computer **43**, 63–70 (2010).
10. J. van der Ham, P. Grosso, F. Dijkstra, A. Taal, and C. de Laat, "On the impact of network topology aggregation in multi-domain lightpath provisioning," Tech. Rep. UVA-SNE-2010-01, Universiteit van Amsterdam (2010).