

Lightpath Planning and Monitoring

Ronald van der Pol¹, Andree Toonk²

¹*SARA, Kruislaan 415, Amsterdam, 1098 SJ, The Netherlands
Tel: +31205928000, Fax: +31206683167, Email: rvd@p@sara.nl*

²*SARA, Kruislaan 415, Amsterdam, 109 SJ., The Netherlands
Tel: +31205928000, Fax: 31206683167, Email: andree@sara.nl*

Abstract: This paper describes the lightpath planning and monitoring software suite developed specifically for managing lightpaths in the hybrid SURFnet6 network and the NetherLight optical exchange point. Lightpaths are high bandwidth circuits with deterministic quality of service characteristics. These lightpaths are used by researchers in demanding network projects. The concept of the high bandwidth lightpaths is new and therefore there is not much network management software available yet that specifically targets these lightpaths. The software suite described in this paper is used for the planning of lightpath resources by automatically finding and reserving paths between two endpoints. The monitoring part of the software suite is used for trouble shooting problems with lightpaths. The status of each lightpath is reported and in case of a problem more detailed information is available to help in identifying the root cause.

1. Introduction

Many National Research and Educational Networks (NRENs) have started to deploy hybrid networks in recent years. These hybrid networks consist of a traditional packet switched (routed) Internet Protocol part (this is the Internet as most of us know it today) and a new and innovative circuit switched optical networking part. These optical circuits are called lightpaths, which are dedicated high bandwidth circuits with capacities of typically 1 to 10 Gbps. As these lightpaths are a new feature in networking, there is not much network management software available yet to plan, provision and monitor these circuits.

This paper describes an open source based network management suite that was specifically developed in order to plan and monitor lightpaths in hybrid networks. This paper starts with an explanation of hybrid networks and the concept of lightpaths. Next, the lack of existing management software for lightpaths and the requirements for a management suite are described.

Finally, the architecture and functionality of the software that was developed is described.

2. Hybrid networks

Hybrid networks offer both routed IP services as well as lightpath services. Lightpaths are point to point connections and are characterized by a very high capacity, a deterministic behaviour and well defined Quality of Service (QoS) parameters. They are typically implemented as TDM circuits or as DWDM wavelengths.

Lightpaths are used for research projects which require high bandwidth services. Applications that use these lightpaths are e.g. large data storage grids like the Large Hadron Collider project at CERN^[1] in Geneva and eVLBI^[2] projects where several radio telescopes are interconnected to form one big virtual telescope.

SURFnet6

SURFnet^[3] is the Dutch national research and education network organisation and was one of the first to build a hybrid network. The SURFnet6 network, is a hybrid network offering both IP and lightpath services. This network was taken into operation at the beginning of 2006.

The IP core is built with highly scalable Avici routers. The SURFnet6 IP network offers both IPv4 and IPV6 unicast and multicast Internet connectivity. The IP network is built on top of the optical network layer.

This optical infrastructure consists of Nortel equipment and 6000 km fiber on which SURFnet has the long-term rights of use. SURFnet6 was developed in the GigaPort project, which runs from 2004 until 2008.

NetherLight and GLIF

The GLIF, Global Lambda Integrated Facility^[4] is an international virtual organization of research networks, research consortia and institutions whose aim is to build a worldwide networking facility for scientific research. The GLIF network is build by connecting individual GOLEs (GLIF Open Lightpath Exchanges) together. Research networks connect to one or more of the GOLEs, resulting in a world wide network of research networks connected to each other by lightpaths.

NetherLight^[5] is one of the major GOLEs in the GLIF network. NetherLight, an initiative of SURFnet, is a distributed facility with points of presence in both Amsterdam and Geneva. Currently over 100Gbps of lightpath capacity is available from and to the NetherLight facility. About half of this capacity is transatlantic.

3. Network management in optical networks

The lightpaths offered in hybrid networks are a new way of network usage. The fact that currently only NRENs offer this capability in their networks is an indication of the novelty of the concept of lightpaths. This also means that the way to implement, manage and operate these lightpaths is not yet clear and is still the subject of much network research. There is also a lot of research being done on how to setup lightpaths that cross several administrative domains. This requires a good understanding of what is requested and what is offered and questions like who is allowed to setup these lightpaths across domains.

Most of the network management tools available today are commercial products developed by the hardware vendors. These products are typically only used to manage the network elements of that specific vendor. Most the products focus on alarm reporting, gathering of performance data and the provisioning of network services such as lightpaths.

The above functionality is what typically is available today in most network management products. However there seems to be minimal interaction between those individual functionalities. In daily operation task, as well as in reporting it is necessary to correlate the lightpath configuration data with for example the alarm data. By doing this it is possible to determine the status of the configured services in the network, and as a result enables the possibility to give accurate reporting as required by Service Level Agreements (SLA). Also lightpath planning and reservations are not something which most of the network management tools provide. Especially the fact that path finding requires multiple constraints from multiple layers, i.e. DWDM and SDH/SONET is something that is not offered by the available products.

In order to operate lightpaths several management tools are needed. For setting up lightpaths a planning and reservation system is needed. This is necessary because often lightpath resources need to be reserved before the actual lightpath is provisioned. For operating and fault management a lightpath monitoring system is needed which provides information about the path taken by the lightpath and information about alarm events. This leads to several requirements for the lightpath management software:

- Reservation of lightpath resources
- Periodic reporting of free resources
- Overview of the status all lightpaths
- Mapping between lightpath and fiber spans (used to determine impact of planned work and fiber cuts)
- Auto configuration, auto discovery
- Open, standards based

Given the specifications described above a software suite was developed in order to plan and monitor hybrid network lightpaths. Because of the open source nature of the software suite it is easy to expand it with new ideas or

adapt it to changes in the way lightpaths are operated. This is important because the concept of lightpaths is new and in constant flux. For lightpath operators it is important to have a network management system that can be easily adapted to changes in the way lightpaths are offered and implemented.

4. Technology architecture

The software suite consists of several building blocks, see figure 1. The TL1 Toolkit is used to extract configuration information from the network equipment. This information is stored in a network state database. The information is also used to generate a topology file of the network. These modules are used by the planning and monitoring server and the information is made available via web services. Each of these building blocks will be described in more details below.

TL1 Toolkit

The TL1 Toolkit^[6] is a Perl Module and is an easy interface to TL1 based equipment. It hides the difficult to understand TL1 commands and output from the user by offering a high level API and by interpreting the output and returning it in easy to use Perl data structures. This toolkit is used to retrieve all of the necessary data from the network elements. Periodically configuration data, inventory data and alarm data is gathered.

The TL1 Toolkit currently offers some predefined functions to gather the crossconnects and alarm information from devices from a range of vendors. Besides that it is possible to use the TL1 Toolkit with any TL1 capable device.

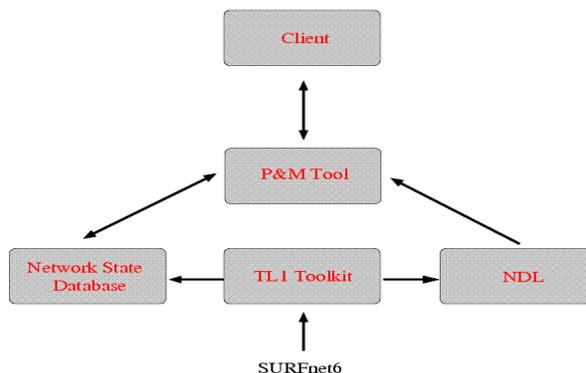


Figure 1 planning and monitoring Software Architecture

Network Description Language (NDL)

The Network Description Language (NDL) framework is formal description of the network and is based on RDF (Resource Description Format), which is a semantic web technique[7]. The NDL file is automatically generated from information retrieved by the network elements by using the TL1 Toolkit and an NDL Perl library[8]. SDH Section Trace is used to determine the neighbours of each network element. The NDL file describes the inventory of the physical network, as well as the topology of the network. Figure 2 is an example of an NDL file.

```
<ndl:Device rdf:about="#Lls001a_ome01">
  <ndl:name>Lls001a_ome01</ndl:name>
  <ndl:locatedAt rdf:resource="#SURFnet6"/>
  <ndl:hasInterface rdf:resource="#Lls001a_ome01:9/1"/>
</ndl:Device>

<ndl:Interface rdf:about="#Lls001a_ome01:9/1">
  <ndl:name>Lls001a_ome01:9/1</ndl:name>
  <ndl:connectedTo rdf:resource="#Zl001a_ome01:5/1"/>
  <ndl:capacity rdf:resource="#OC192"/>
  <ndl:encapsulation rdf:resource="#SONET"/>
</ndl:Interface>

<ndl:Device rdf:about="#Zl001a_ome01">
  <ndl:name>Zl001a_ome01</ndl:name>
  <ndl:locatedAt rdf:resource="#SURFnet6"/>
  <ndl:hasInterface rdf:resource="#Zl001a_ome01:1/1"/>
  <ndl:hasInterface rdf:resource="#Zl001a_ome01:5/1"/>
</ndl:Device>

<ndl:Interface rdf:about="#Zl001a_ome01:5/1">
  <ndl:name>Zl001a_ome01:5/1</ndl:name>
  <ndl:connectedTo rdf:resource="#Lls001a_ome01:9/1"/>
  <ndl:capacity rdf:resource="#OC192"/>
  <ndl:encapsulation rdf:resource="#SONET"/>
</ndl:Interface>
```

Figure 2: Example NDL description

Network state database

The network state database contains all the dynamic information from the network. Network state information, like crossconnects and alarm status information from the network is frequently and automatically synchronized with the network state database.

A lightpath consists of multiple crossconnects on multiple network elements. A crossconnects consist of an identifier, an incoming and outgoing interface (and a switchmate for protected lightpaths) and the associated timeslot information. Lightpath reservations are also stored in the network state database. Alarm information like the type of alarm, on which network elements and the impacted interfaces are stored in the same database. This enables the possibility to correlate the different sources of information.

5. Lightpath planning

An important requirement for this network management suite was to be able to do lightpath planning and reservations. The chapters above describe the necessary sources of information, i.e. an NDL file describing the network inventory and topology as well as the network state database which holds the current state of the network. The lightpath planning tool is used by network operators who are responsible for planning and reserving network resources. The planning tool consists of a path finding module where an optimal path through the network is found depending on several constraints. The resources to be used for the path can then be stored in the network state database and flagged as reserved.

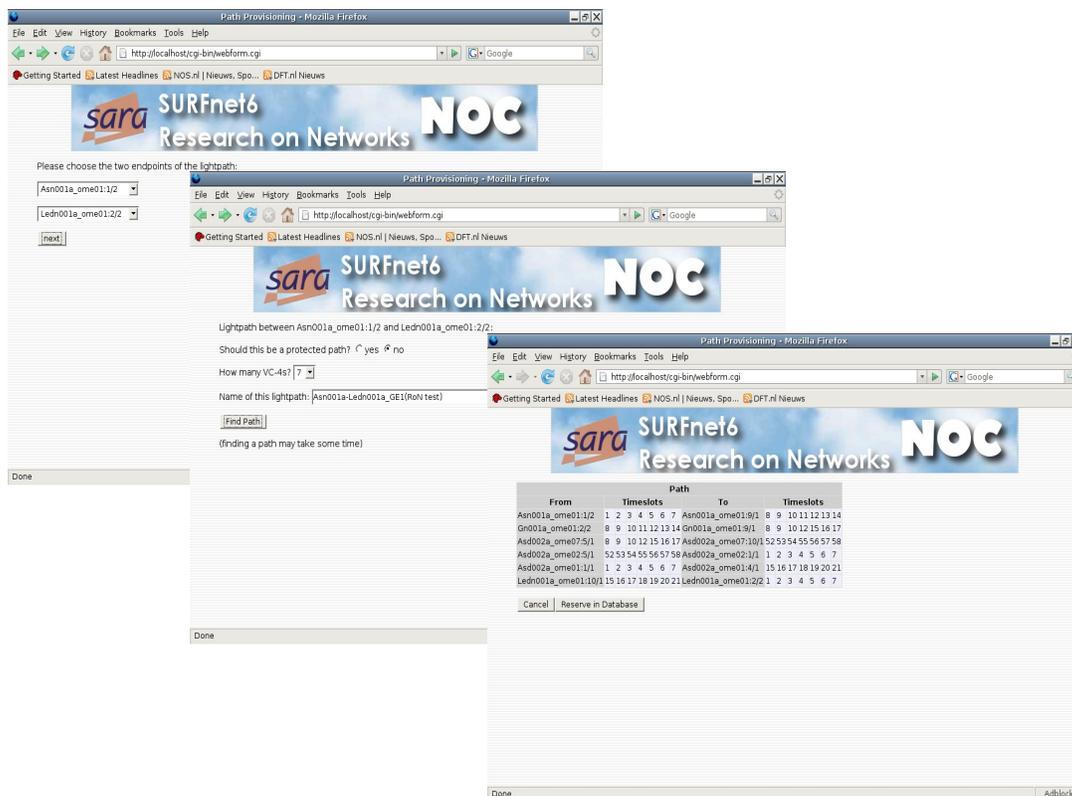


Figure 3: Lightpath planning in SURFnet6

The planning tool, as displayed in figure 3, offers a web interface to the user. First the user is asked to select two end-points from a list of all the available end-points. This list is created by selecting all the end-point interfaces from the NDL file, and checking with the network state database for availability. Only the available interfaces are presented to the user. The next step is to select the capacity and whether this should be a protected or unprotected path. Once this information is known by the application it will start the path finding

process. First a graph of the network is constructed with the topology information as described in the NDL file. Then all the links, the edges in the graph, which do not meet the capacity constraints, are pruned from the graph. This is done by querying the network state database. Now a graph is available which meets the capacity constraints of the request, this graph is used as input for the path finding algorithm. A constrained based Dijkstra shortest path first algorithm is used to find a path between the two selected endpoints. The current constraints are link usage and cost. Some links are more preferred than others, this is reflected in a cost value in the graph. Another constraint is the amount of available free timeslots on each link. Links with many available timeslots are preferred above links with few available timeslots. If the request was to find a protected path, then the links and interfaces used by the primary path are removed from the graph. Then the constrained based Dijkstra shortest path first algorithm is used again with the modified graph to find a backup path.

In the next screen the user is presented with the results. A list of network elements is given in topological order together with the associated timeslot information; this represents the path(s) found. If the user is satisfied with the result, the path can be reserved in the network state database.

6. Lightpath monitoring

Lightpath monitoring is an important aspect in the daily network management tasks. With the currently used commercial product the NOC will see alarms on the network, however these are not correlated to lightpath services. The network state database as described in this paper contains both the details of the configured network services as well as the current alarm information in the network. This allows the application to first identify all the unique lightpaths in the network, this is done by combining all the relevant crossconnects together. The next step is to correlate the interfaces and associate timeslots of the lightpaths, with the current alarm information.

The result^[9] is a list of lightpaths with corresponding alarm information. This list gives an overview of all the lightpaths with a status of up or down, see figure 4. Because this list is generated, new lightpaths appear automatically without the need to edit any configuration files. This gives the user in one view the current status of the lightpaths. It is also possible to click on one of the lightpaths, which will show a more detailed page. This includes the configuration details, alarm information, as well as a visual representation of the lightpath.

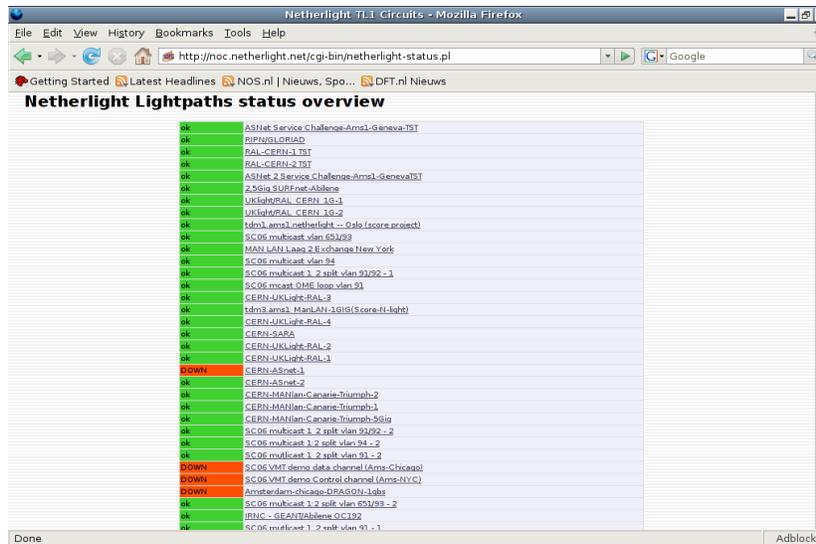


Figure 4 Lightpath monitoring in NetherLight

7. Conclusions and future work

Lightpaths are a new concept and therefore there is not much out of the box software available yet. So we built our own planning and monitoring suite based on open software and standards. We are able to automatically find paths through the network and store these as reserved lightpaths in a network state database. For monitoring a website shows an overview of all the lightpaths with associated alarm status information. All information is automatically retrieved from the network elements, so there is no need to add new lightpaths to a configuration file. Newly provisioned lightpaths appear automatically in the list of all lightpaths. This software is in daily use by a European NREN and also by a large optical exchange point in Europe.

The lightpath planning and monitoring software suite will be developed further. Currently the planning tool it is being extended with Shared Risk Link Group (SRLG) support. This means that also the DWDM layer will be included in the network graph which is used for path finding. Protected paths will get the additional requirement that the primary and backup path should not have a SRLG in common.

Work will continue in order to support interdomain planning and monitoring. Currently we are analyzing what would be needed to support true interdomain monitoring as well as path finding. To realize that domains would somehow need to publish lightpath status and/or alarm information as well as topology information. Parts of this have already been implemented and tested. At SuperComputing 2006 interdomain path finding with the help of NDL was demonstrated [10].

The visualisation of lightpaths will also be enhanced in future versions. Especially visualizing complex graphs, both intra and interdomain as well as different protection schemes will need to be supported.

8. References

- 1[] CERN Large Hydron Collider: <http://lh.web.cern.ch/lhc/>
- 2[] The European VLBI Network : <http://www.evlbi.org/>
- 3[] SURFnet: <http://www.surfnet.nl>
- 4[] GLIF, the Global Lambda Integrated Facility: <http://www.glif.is>
- 5[] NetherLight: <http://www.netherlight.net>
- 6[] TL1 Toolkit: <https://noc.sara.nl/nrg/TL1-Toolkit/>
- 7[] Jeroen van der Ham, Paola Grosso, Ronald van der Pol, Andree Toonk and Cees de Laat. Using the Network Description Language in Optical Networks. , May, 2006, Accepted by the IEEE Integrated Management Conference 2007: <http://staff.science.uva.nl/~vdham/research/publications/0606-UsingNDLInOpticalNetworks.pdf>
- 8[] NDL Perl libraries: <https://noc.sara.nl/nrg/ndl/>
- 9[] NetherLight lightpath status overview: <http://noc.netherlight.net/cgi-bin/netherlight-status.pl>
- 10[] SuperComputing2006 demonstration, Interdomain path finding with NDL: <http://traffilight.uva.netherlight.nl/NDL-demo/sc-demo-test/pathfind-new.htm>