

Lightpath Planning and Monitoring in SURFnet6 and NetherLight

Ronald van der Pol

SARA, Kruislaan 415, 1098 SJ Amsterdam, The Netherlands
e-mail: rvd@p@sara.nl

Andree Toonk

SARA, Kruislaan 415, 1098 SJ Amsterdam, The Netherlands
e-mail: andree@sara.nl

1 Abstract

SURFnet is operating a hybrid network with both IP and lightpath services. The lightpath service is new and there is not much software available yet for managing this service. These lightpaths are also available on NetherLight, Europe's largest Optical Exchange located in Amsterdam which is operated by SURFnet. This paper describes the open standards based software suite that was written to do the planning and monitoring of lightpaths. The software allows the NOC to find optimal paths between two points in the network and to make lightpath reservations. It also describes several monitoring tools that are used in both the SURFnet6 hybrid network and the NetherLight Optical Exchange.

2 Keywords

lightpath, monitoring, optical networking, path finding, hybrid networks

3 Introduction

Many National Research and Educational Networks (NRENs) have started to deploy *hybrid networks*. These networks consist of a traditional routed IP part and an optical circuit-switched part with *lightpath* services. Lightpaths are dedicated high-speed point-to-point circuits with a capacity of typically one to ten Gbps and with well-defined Quality of Service (QoS) parameters. These lightpaths can be setup between two points within the NREN, but also world wide between different NRENs. This is done by using Optical Exchange Points like NetherLight.

3.1 SURFnet6

The hybrid SURFnet6 network is operational since the beginning of 2006. It offers IPv4 and IPv6 unicast and multicast as well as lightpath services. In order to be able to deliver these lightpath services SURFnet has procured 6000 kilometre of dark fiber and installed their own DWDM and SDH equipment on it. The DWDM layer consists of Nortel's Common Photonic Layer (CPL) which offers up to 72 wavelengths (lambdas) per fiber. As each wavelength can operate at 10 Gbps, a fiber can transport up to 720 Gbps. Some wavelengths are operated as 10 Gbps lightpaths, but most of them are used by the SDH layer and subdivided in (typically) 1 Gbps lightpaths. The Nortel OME6500 equipment is used to provide customers with 1GE or 10GE Ethernet connections. Lightpaths can be offered as unprotected or as protected circuits.

3.2 NetherLight

NetherLight is one of the largest Optical Exchanges in the world. An Optical Exchange Point is a concentration point where organisations can connect their optical transmission lines (typically OC-192 lines) in order to accommodate lightpaths across several domains. NetherLight as a major exchange point plays an important role in the world wide optical networks. NetherLight participates in several international projects:

- NetherLight is a participant in the Global Lambda Integrated Facility (GLIF). The GLIF participants jointly make lambdas available as an integrated global facility for use by scientists and projects involved in data-intensive scientific research.

- NetherLight is also a partner in GLORIAD (Global Ring Network for Advanced Applications and Development), which has built an optical ring around the globe.

4 Network Management in Optical Networks

Hybrid networks have opened a new territory for most NRENs as they are now also responsible for planning and monitoring the DWDM infrastructure and the lightpath services on it. Existing network management tools have several shortcomings. Typically, they monitor individual network equipment only. Equipment failure and interface outages are shown, but usually there is no link to the services that are affected by the fault. However, this is exactly what is needed for lightpath monitoring. A fault in the network has to be correlated to the lightpaths. Ideally, this correlation is automatically retrieved from the network. In order to accomplish this, for each lightpath the network elements and interfaces used must be determined automatically.

This has led to the following requirements for the lightpath monitoring software:

- overview of the status all lightpaths
- mapping between lightpath and fiber spans (used to determine impact of planned work and fiber cuts)
- autoconfiguration, autodiscovery
- open, standards based

When planning a new lightpath the topology of the network as well as information about the available timeslots is needed. Usually, all links in the network will also have a metric to prefer certain links above other links. A constrained based shortest path first algorithm can then be used to find a path through the network.

5 Architecture

The software suite consists of several building blocks (see figure 1). SARA's TL1 Toolkit is used to extract the current configuration information from the network equipment. This information is stored in the Network State Database. The information that is retrieved from the network elements is also used to automatically generate a description of the network. These modules are used by the Planning and Monitoring Server and the information is made available via web services.

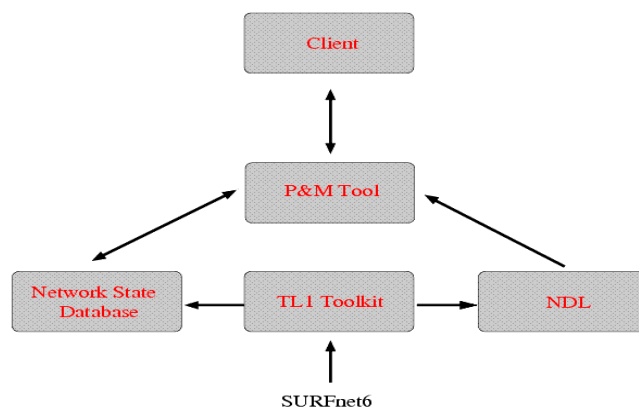


Figure 1 Planning and Monitoring Software Architecture

5.1 TL1 Toolkit

The TL1 Toolkit [1] is a Perl Module written by SARA and provides an easy interface to TL1 based equipment. It hides the sometimes complex and difficult to understand TL1 commands and output from the user. The TL1 Toolkit offers a high level Perl API, which allows a user to very easy execute TL1 commands on any TL1 capable device. The TL1

Toolkit will interpret the TL1 output and return this output in a Perl data structures. This allows a network operator to write different kinds of scripts and make use of the TL1 capabilities of the network device without having to worry about the exact syntax.

The TL1 Toolkit offers a set of pre-defined functions which can be used to for example retrieve all the current alarms of a network element, or to retrieve all the crossconnects on a network element. These functions support different types of equipment from different vendors. The output is returned in a standardized data structure, this means that the user does not need to worry about the different forms of output and parsing logic. Currently there are several predefined functions for Nortel CPL (DWDM), Nortel OME6500, Nortel HDXc and the Cisco ONS15454. The TL1 Toolkit has also been tested with Nortel OM5200 and Adva equipment. Multiple demonstration applications have been written such as a TL1 MRTG plugin, a TL1 to SNMP trap script and a TL1 nagios plugin.

5.2 Network Description Language

The Network Description Language (NDL) framework is being developed by the University of Amsterdam [2]. It is a formal description of a network and is based on RDF (Resource Description Format), which is a semantic web technique. The planning and monitoring software uses an NDL file as input to determine the topology and properties of the network. The NDL file describes which network elements are in the network, what kind of interfaces this network element has and some specific properties about the interfaces. This information is used by the planning and monitoring software to get a detailed description of the network and also to construct the network graph.

The NDL file is generated from information that is retrieved from the network elements by using the TL1 Toolkit and an NDL Perl library [3]. The NDL generator script retrieves information such as the inventory and the SDH Section Trace. This is used to determine the neighbours of an interface in the network element. Section Trace is a feature in the SDH and SONET protocol stack, which allows you to send a string in the overhead bytes. This string is used to represent the neighbouring interface. Once all this information is available it can be combined and eventually be used to create an NDL file. Currently the SURFnet6 NDL file consists of approximately 5000 lines of generated NDL lines.

In figure 2, part of an NDL file is shown. The device part describes the device name and which interfaces are in the device. Next there is the description of all the interfaces in the device. This contains the encapsulation and the capacity of the interface. If the encapsulation type is *sonet*, then there is also a *connectedTo* definition that describes the neighbour.

```
<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="#Z1001a_ome01:5/1"/>
  <ndl:capacity rdf:resource="#OC192"/>
  <ndl:encapsulation rdf:resource="#SONET"/>
</ndl:Interface>

<ndl:Device rdf:about="#Z1001a_ome01">
  <ndl:name>Z1001a_ome01</ndl:name>
  <ndl:locatedAt rdf:resource="#SURFnet6"/>
  <ndl:hasInterface rdf:resource="#Z1001a_ome01:1/1"/>
  <ndl:hasInterface rdf:resource="#Z1001a_ome01:5/1"/>
</ndl:Device>

<ndl:Interface rdf:about="#Z1001a_ome01:5/1">
  <ndl:name>Z1001a_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

5.3 Network State Database

The Network State Database is used to store information about the network, like crossconnects and alarm information. A crossconnect consists of an incoming and outgoing interface (and a switchmate for protected lightpaths) as well as the associated timeslot information. This information is periodically retrieved from the network elements using the TL1 Toolkit and is stored in the Network State Database. The same principal is used for alarm information. This is necessary in order to determine the operational state of the network which is used by the monitoring software. Lightpath reservations are also stored in the Network State Database.

The Network State Database also contains state information about the DWDM layer. Information about how different wavelengths are routed over the all-optical backbone is periodically retrieved and stored in the database. This information provides the necessary data for generating Shared Risk Link Groups (SRLG). This is important information for planning protected lightpaths.

The Network State Database is currently implemented by using a MySQL database.

6 Lightpath planning in SURFnet6

The planning tool consists of a path finding module which will try to find an optimal path through the network based on several constraints. The resources to be used for the path can then be stored in the Network State Database and flagged as reserved. The planning tool offers a web interface to the user (see figure 3). The user is asked for two end-points, the capacity and whether protection is needed. As a result a list of network elements is given in topological order together with timeslot information.

The NDL topology file contains information about devices and interfaces and which interfaces are connected to each other. This information is read by the planning tool. The topology information is stored in a graph data structure. The vertices of the graph are the device-interface combinations and the edges of the graph are the links between the interfaces. The NDL file also stores information about the type and capacity of the interfaces. The 1GE interfaces with no neighbour are the edge interfaces where customers can connect to. A list of these interfaces is presented to the user to choose as end-points of the lightpath.

The Network State Database contains information about timeslots used on all the interfaces. This information is used to prune those edges from the graph that do not meet the capacity constraints for the requested lightpath. Each edge is also given a metric. This is used to prefer certain paths in the network above others. These metrics are based on several constraints. One of them is the amount of available timeslots. Links with many available timeslots are preferred above links with few available timeslots.

A constrained based Dijkstra shortest path first algorithm is used to find a path between two end-ports in the network. When the user requests a protected path, the Dijkstra algorithm is run twice. After the first shortest path is found, the network elements and interfaces used by that path are pruned from the graph. In this modified graph the Dijkstra algorithm is run for a second time to find a backup path that does not use network elements or interfaces from the primary path. It is known that this does not always result in a optimal set of paths. A future improvement to the algorithm would be to use a shortest cycle algorithm.

When the protected or unprotected path is found and the user is satisfied with the results, the interfaces and timeslots used by the path can be reserved in the Network State Database.

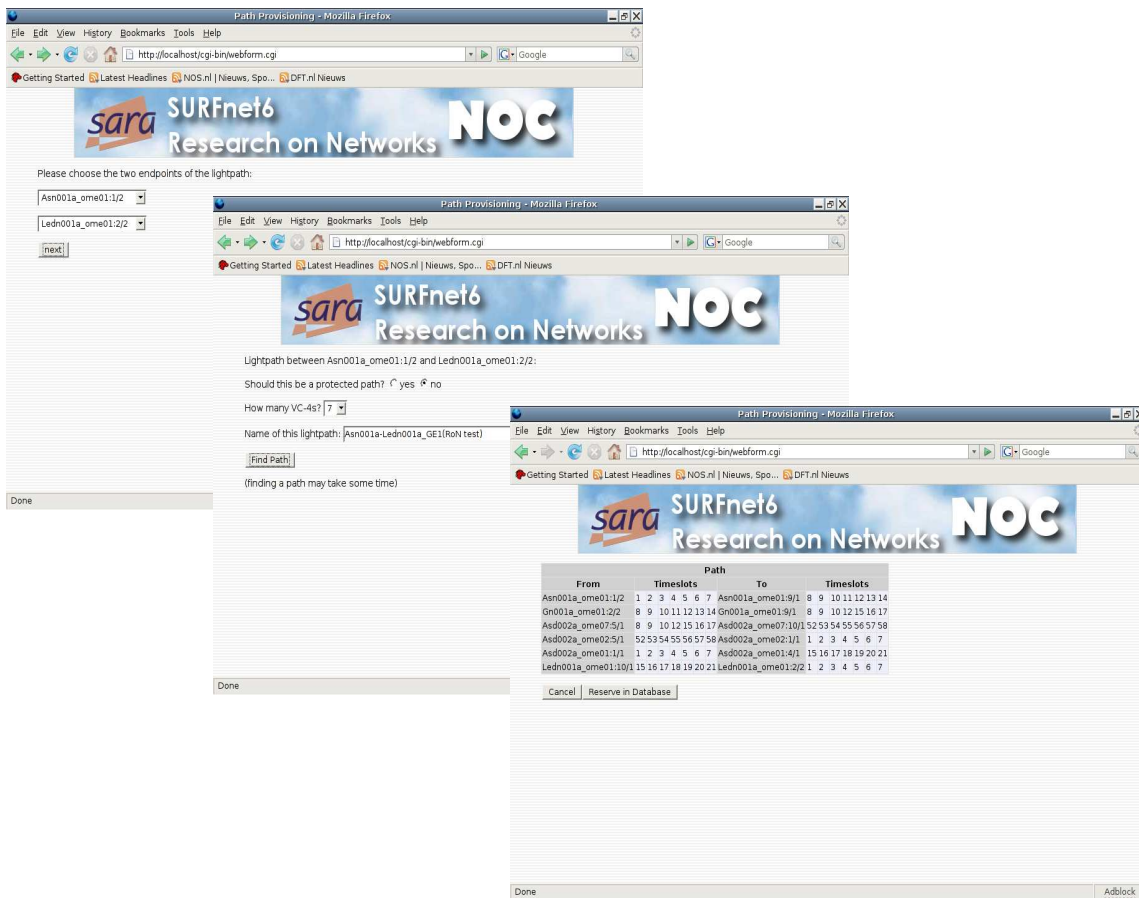


Figure 3 Lightpath planning in SURFnet6

7 Lightpath monitoring in NetherLight

The TL1 Toolkit is used to periodically retrieve alarm and crossconnect information from the network elements. The alarms are related to specific interfaces and timeslots. The crossconnect information with the Circuit ID gives information about the devices and interfaces used by the lightpaths. All this information is stored in the Network State Database. This database also contains information about all the lightpaths.

The alarms are then correlated to the lightpaths and as a result a list of lightpaths is generated with corresponding status information. This list gives an overview of all the lightpaths with a status of *up* or *down*. See figure 4 and [4]. Because this list is generated by periodically retrieving information from the network elements, new lightpaths appear automatically without the need to edit any configuration files.

The list of lightpaths is presented as a web page with the Circuit IDs as hyperlinks. These hyperlinks point to detailed information about the lightpath. In this detailed view a list of all the network elements is shown, the interfaces that are used and the timeslots used on those interfaces. It also shows on which interface an alarm is present and a description of the alarm.

A graphical presentation of the lightpath is given showing the devices and interfaces that the lightpath uses. The links that are down are drawn in red.

This monitoring service was setup as an aid to fault resolution. The overview page with all the lightpaths is a quick way to identify which lightpaths are down. Clicking on the lightpath gives more detailed information about the section that is probably affected. In this way the engineer can quickly zoom into the problem spot.

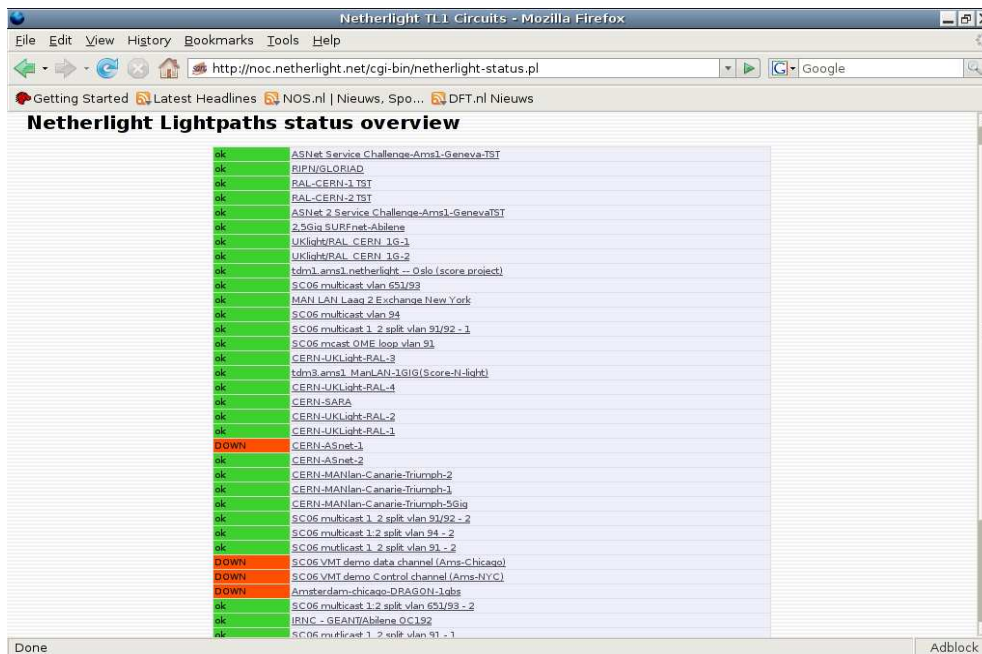


Figure 4 Lightpath monitoring in NetherLight

8 Future Work

The lightpath planning and monitoring software suite will be developed further. The planning part will be enhanced with Shared Risk Link Group (SRLG) support. SRLGs are different links that are known to fail at the same moment when there is e.g. a fiber cut. Examples are fibers that use the same duct. The planning tool will get a way to associate a SRLG with each link. Protected paths will get the additional requirement that the primary and backup path may not have a SRLG in common.

The monitoring tool will be enhanced to also support lightpaths that cross several domains. This will be done by using lightpath information from other domains. This can be done by sharing NDL files between domains or by offering the information via web services or both. The web page will show the true end-to-end lightpath information and will show in which domain and section the problem is. This will be a great help in the inter-domain fault resolution of lightpaths.

9 Conclusion

Lightpaths are a new concept and therefore there is not much out of the box software available yet. So we build 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.

References

- [1] <http://nrg.sara.nl/TL1-Toolkit/>
- [2] Using the Network Description Language in Optical Networks, *Jeroen van der Ham, Paola Grosso, Ronald van der Pol, Andree Toonk, Cees de Laat*, accepted for publication by IM 2007
- [3] <http://nrg.sara.nl/ndl>
- [4] <http://noc.netherlight.net/>