

# Project Report RP2

*The PerfSONAR framework for inter-domain monitoring in the  
GLIF infrastructure*



## **Authors:**

ing. Tijmen van den Brink  
ing. Peter Tavenier

## **Supervisors:**

R. van der Pol, MSc (SARA)  
A. Toonk, MSc (SARA)  
dr. P. Grosso (UvA - SNE research group)

July 5, 2007

## **Abstract**

PerfSONAR is a network monitoring framework based on a service oriented architecture. This projects primary objective was to research the possible use of the PerfSONAR framework for inter-domain monitoring in the GLIF infrastructure. This, including a comparison of the functionality of PerfSONAR and Spotlight (a monitoring tool developed by SARA) and the advantages that can be gained by implementing the PerfSONAR framework possibly in combination with Spotlight.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Problem description</b>	<b>4</b>
<b>3</b>	<b>Related work</b>	<b>5</b>
<b>4</b>	<b>Global Lambda Integrated Facility</b>	<b>6</b>
4.1	About GLIF . . . . .	6
4.2	Requirements . . . . .	6
4.3	Functional requirements . . . . .	8
<b>5</b>	<b>PerfSONAR framework</b>	<b>9</b>
5.1	PerfSONAR architecture . . . . .	9
5.2	PerfSONAR services . . . . .	10
5.3	Communication between services . . . . .	12
<b>6</b>	<b>E2E inter-domain monitoring in the GLIF infrastructure</b>	<b>13</b>
6.1	E2E Monitoring System . . . . .	13
6.2	Concepts . . . . .	13
6.3	Data provisioning . . . . .	16
6.4	PerfSONAR and Spotlight . . . . .	17
<b>7</b>	<b>Future work</b>	<b>18</b>
<b>8</b>	<b>Conclusions</b>	<b>18</b>
	<b>Acknowledgments</b>	<b>18</b>
	<b>References</b>	<b>19</b>
	<b>Appendices</b>	<b>21</b>
<b>A</b>	<b>XML E2E-link configuration files</b>	<b>21</b>
A.1	SURFnet domain . . . . .	21
A.2	GEANT2 domain . . . . .	22
A.3	CERN domain . . . . .	26

# 1 Introduction

From June 4 until July 4, 2007, Peter Tavenier and Tijmen van den Brink did a research project on PerfSONAR. This project was initiated by the “Universiteit van Amsterdam (UvA)” and supported by SARA (R. vd Pol, MSc and A. Toonk, MSc) and the UvA SNE research group (dr. P. Grosso). PerfSONAR is a network monitoring framework developed by a joint collaboration of ESnet, GEANT2, Internet2 and RNP. The main purpose of this research is to find out if PerfSONAR is suitable for inter-domain lightpath monitoring in the GLIF infrastructure.

The rest of this report is organized as follows. In section 2 the research question will be clarified with a problem description. Related work concerning this research will be presented in section 3. This research is focussed on the GLIF infrastructure. GLIF and their requirements will be highlighted in section 4. PerfSONAR consist of different services. These services and the PerfSONAR framework will be presented in section 5. A specific PerfSONAR monitor system, E2E monitoring system, which can be used for monitoring E2E links will be explained in section 6. The conclusions and future work are given in the last two sections.

## 2 Problem description

Many scientific research projects create big amounts of data that need to be transported worldwide. Today lightpaths are used to transfer this data between scientific institutes. A lightpath is a point-to-point, circuit switched optical path, between two nodes with a deterministic character. This means guaranteed Quality of Service (QoS). Lightpaths often run through multiple domains because of the worldwide collaboration of research institutes.

Monitoring a lightpath in a single domain is done by a NOC (Network Operation Control) of that specific domain. However monitoring a lightpath crossing multiple domains (inter-domain lightpath, see figure 1) is different. Every NOC can see a part of the lightpath, yet a total overview is missing. To get a total overview of the status of a lightpath all the NOCs need to join their information. Viewing this combined information with a client gives you an overview of the total status of the lightpath.

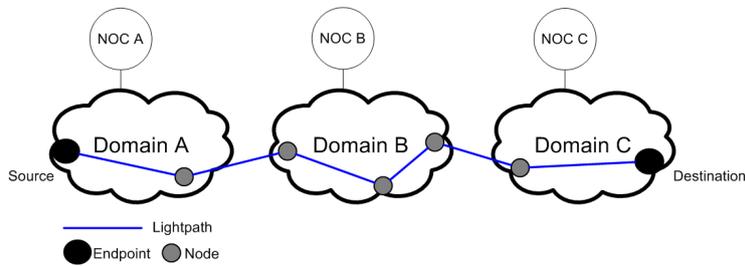


Figure 1: Inter-domain lightpath

PerfSONAR[14] is a network monitoring framework based on a service oriented architecture. Implementing this framework in the different domains should make it possible to monitor an inter-domain lightpath. Therefore the primary objective of this research is to determine if it is possible to use PerfSONAR for inter-domain lightpath monitoring. An example of an inter-domain network is the Global Lambda Integrated Facility infrastructure, described in section 4.

### 3 Related work

The first release of the PerfSONAR project, PerfSONAR v1.0 was July 18, 2006. June 5, 2007 version 2.0 was released. Therefore PerfSONAR described in this report refers to PerfSONAR version 2.0.

Previous work related to the PerfSONAR project are papers written by the participants of PerfSONAR. These papers give a clear explanation of the PerfSONAR framework.

A good theoretical description of the inter-domain monitoring aspects is given in [1]. Practical implementations of this description are given in [5].

In [6] the visualization tools of PerfSONAR are described. This paper gives a good indication of what can be visualized within the PerfSONAR framework and is therefore mentioned here.

The service oriented architecture of PerfSONAR is clearly explained in [7]. It gives a good overview of the communication between the several PerfSONAR services.

This document will refer to several documents which are also related to this project or clarify some specific subjects.

## 4 Global Lambda Integrated Facility

### 4.1 About GLIF

GLIF stands for Global Lambda Integrated Facility. It is a virtual organization with mostly National Research and Educational Networks (NRENs) as participants. As stated in [4] the two main activities of GLIF are:

- To jointly make lambdas<sup>1</sup> (i.e. wavelengths on a single fiber) available as an integrated global facility for use by scientists and projects involved in data-intensive scientific research.
- GLIF brings together leading networking engineers worldwide, who exchange information to learn from each others experiences, seek to establish best practice, work together to develop, test and implement new lambda networking technologies, middleware and applications, and generally collaborate to bring the technology forward.

Furthermore the GLIF community shares a common vision that focuses on optical networks to provide the needs of today's most demanding e-science applications. Lambda networking in GLIF terms is about offering lambdas (i.e. wavelengths) to user communities with guaranteed Quality of Service (QoS) on a shared fiber infrastructure.

### 4.2 Requirements

The GLIF community involves a lot of parties which all have their own requirements. In [16] stakeholders involved in a NREN organization are identified as:

- The general public and policy makers;
- Connected organizations;
- The network operator itself, including other connected NRENs;

The next subsections will discuss the requirements of the different stakeholders.

#### 4.2.1 The general public and policy makers

NRENs are most likely to be funded through public funds. This necessitates the organization to provide some information about the value of the network to justify expenditures [16]. The information provided should give a global view on the usage of the network. A suitable metric is the amount of data per time unit transported through the network. Note that different services (e.g. IP, or at a lower layer Ethernet) should be treated individually. The general public and policy makers require only a simple representation of the network usage.

---

<sup>1</sup>Wavelength Division Multiplexing is a technique to multiplex multiple optical carrier signals on a single fiber by using different wavelengths.

### 4.2.2 Connected organizations

A more detailed set of information should be made available to connected organizations. To show the clients their usage of the network the amount of data per time unit transported through each of the acquired interfaces is suitable. This helps to foresee future bottlenecks, justify the expenditure and abnormal traffic - possibly indicating errors or security incidents - will be noticed[16]. Preferably the client wants to relate the traffic on a certain interface with internal users. Distinguishing traffic on end-point specific characteristics is preferred.

### 4.2.3 Network operators and connected NRENs

As a network operator you want full information of your network. The parties involved in this information or a subset of this information are 1) the stakeholders mentioned in sections 4.2.1 and 4.2.2, 2) the network operator itself and 3) the connected NRENs. The first is provided only a subset of the information. The network operator itself is interested in all of the detailed information that is extracted from the network. This information is useful for several aspects including accounting, charging, capacity planning and network design. The third group - the connected NRENs - introduces the inter-domain aspect. When lightpaths cross different domains (see figure 1) the need for other NRENs to have access to detailed information (e.g. link status and network performance) of involved domains increases.

If a lightpath goes down NOC A can only verify that the problem does not exist in domain A. When all NOCs are sharing their network information the problem could be solved in a more efficient way.

### 4.2.4 Metrics in the GLIF infrastructure

Although the GLIF participants are primarily offering services at the lower layers (i.e. 0 and 1), but also at layer 2. The monitoring requirements will include services at layer 0,1 and 2. Figure 2 depicts the offered service at each layer. At

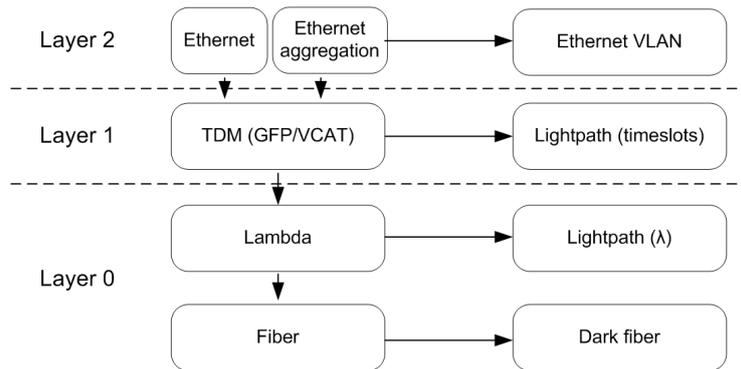


Figure 2: Services and capabilities in the hybrid network [9]

the lowest layer (i.e. layer 0) dark fibers and with the use of WDM techniques lambdas can be provided. At this layer the following metrics/information could be interesting:

- The signal strength/optical loss of the wavelengths
- Link status; up/down - availability

At layer 1 SONET or SDH framing is used to create TDM based paths which divide a single lambda in a number of time slots depending on the protocol. The metrics of interest are:

- Errored Second Count (ES)
- Unavailable Seconds (US)
- Available timeslots
- Link status; up/down - availability

To optimize the use of the layer 2 Ethernet service, VLAN technology can be used to aggregate multiple clients while remaining separate datastreams [16]. Useful metrics are:

- Runts (a packet that is too small)
- Giants (a packet that exceeds the maximum size)
- CRC errors
- Interface errors/drops
- Number of packets sent/received
- Link status; up/down - availability

### 4.3 Functional requirements

One of the functional requirements for the GLIF infrastructure is that the monitoring tool must support unprotected and protected lightpaths. A protected

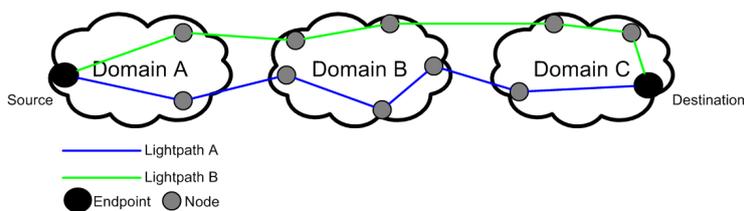


Figure 3: protected lightpath

lightpath connects the same endpoints as the unprotected lightpath via a different route (see figure 3). The connection between the two endpoints remains when a fiber cut takes place in lightpath A, because lightpath B is the backup route. In monitor visualization tools a fiber cut in lightpath A should be visible and also that lightpath B is still up. However in a total overview it should be visible that the lightpath is in operation.

Some lightpaths are only needed for a short time project. When a lightpath is set up, it must be easy to configure, or preferable configured automatically. Also when a lightpath is removed the configuration should be easy to maintain.

Another functional requirement is providing authentication and authorization functionality.

## 5 PerfSONAR framework

PerfSONAR stands for PERFORMANCE focused Service Oriented Network monitoring ARCHitecture. It is a joint collaboration between ESnet[2], GEANT2[3], Internet2[10] and RNP[15]. PerfSONAR is a monitoring framework based on a service oriented architecture (SOA)[13]. This type of architecture has proven to be very flexible, robust and scalable. This makes it suitable for multi-domain Network monitoring[7].

### 5.1 PerfSONAR architecture

It is possible to divide the PerfSONAR architecture into three layers. As stated in [1] and depicted in figure 4 the three layers are:

- User interface layer
- Service layer
- Measurement point layer

Each of these layers are described in detail in the subsections 5.1.1 to 5.1.3.

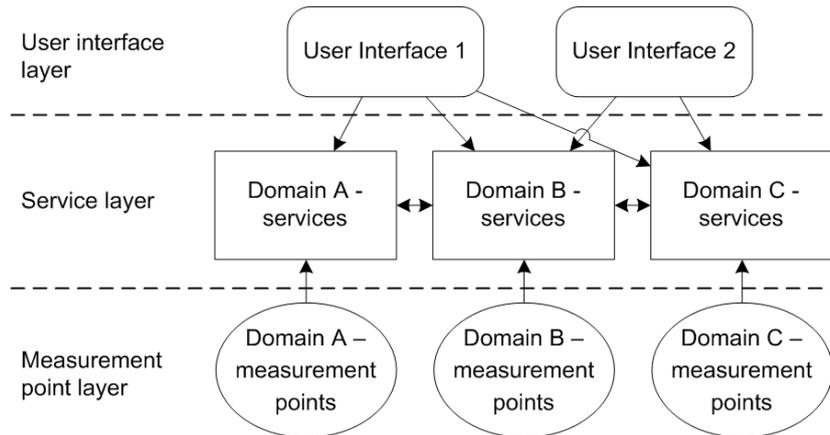


Figure 4: Three layer model in the PerfSONAR architecture[1]

#### 5.1.1 Measurement point layer

At this layer all the monitoring takes place. The measurement point service monitors the network and provides the network information to the upper layers. Monitoring can be arranged in active monitoring and passive monitoring.

Passive monitoring is a way of monitoring where the statistics are retrieved from the devices (e.g. SNMP to retrieve statistics).

Active monitoring is performing tests to get your performance statistics out of the network. An example tool is Bandwidth test controller (BWCTL[11]) which uses Iperf[12] to perform active monitoring. For example Iperf generates a huge amount of traffic to transfer over a link to determine the maximum throughput of that link.

### 5.1.2 Service layer

The service layer is the middle layer in the PerfSONAR framework. The main purpose is to interact between the user interface layer and the measurement point layer and it is therefore responsible for the domain control. This is done by following services; lookup service, measurement archive service, authentication service, transformation service and resource protector service. Some of these services contain an interface and can be contacted directly by the user interface layer.

In figure 4 the lines that connect the different domain at the service layer are indicating the information exchange of the lookup services (explained later on in section 5.2.1). The other services in the service layer interact with services in their own domain or with the user interface of another domain.

### 5.1.3 User interface layer

As the name implies, the user interface layer is an interface for the user to the lower layers. In this layer the visualization tools (e.g. analysis tools, threshold alarms, visualization graphs) show what is in the layers below. Adapting tools to the user specific needs is possible so that many different types of users get their information out of the framework. The user interface layer adds a level of abstraction by concealing all the measurement points and services in the lower layers.

The user interface first contacts the service layer at a lookup service (section 5.2.1) to request which measurement data is available and where it can be found. Therefore the interface must know the address of the lookup service. This address can be found in another lookup service or must be “hard-coded”.

## 5.2 PerfSONAR services

PerfSONAR consists of several services, because of its service oriented architecture. The services described in the PerfSONAR online documentation are:

- Lookup Service
- Measurement Point Service
- Measurement Archive Service
- Authentication Service
- Transformation Service
- Resource Protector Service

Each service is able to run independently, however it is also able to interact with other services. This communication between services is based on SOAP XML messages and uses the schema of the Global Grid Forum (GGF) Network Measurement Working Group (NMWG). This NMWG schema is briefly explained in section 5.3 and described in detail in [17].

Figure 5 displays which services contact which services and for what purpose. These services are also explained separately in subsections 5.2.1 to 5.2.3. Each service can be implemented multiple times.

The services for authentication, transformation and resource protection, described in subsections 5.2.4 to 5.2.6, are only described in the documentation as how they should work. There are not yet implementations available and these services are not used in this project. Therefore they are not mentioned in figure 5.

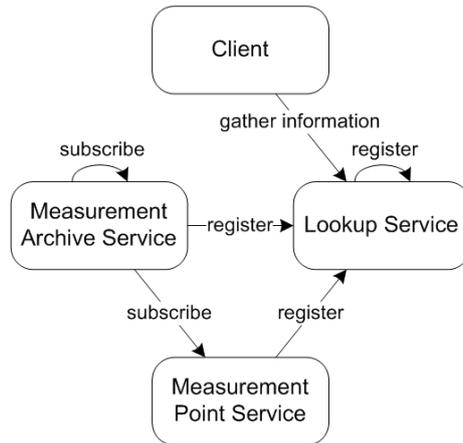


Figure 5: Interaction between services within the PerfSONAR framework[7]

### 5.2.1 Lookup service

The lookup service acts as a directory service. Services register themselves and their capabilities with the lookup service (see figure 5). A lookup service can also register itself to other lookup services. However in a typical situation every domain has only one lookup service.

The client (user interface) communicates with this service to get information about the other services. Clients are able to query the lookup service to acquire the available services. Therefore the client must have knowledge of the address of a lookup service. Queries to the lookup service are based on attributes like: service type, required authentication attributes, service capabilities, network location, etc.

The information is stored in a XML database with XQuery capabilities[19]. XQuery is a query language that is designed to query collections of XML data.

### 5.2.2 Measurement point service

A measurement point service provides the monitoring information. PerfSONAR has several implementations of measurements points, i.e. Bandwidth test controller (BWCTL), One-Way Active Measurement Protocol (OWAMP) and End-to-End Monitoring (E2EMon). These measurement services are rolled out at the moment as micro-releases, this means that they are not fully implemented in the perfsonar bundle yet.

The measurement point service pushes information to all subscribed measurement archives (section 5.2.3). This information is in compliance with the NMWG XML schema format.

It is possible to use any monitoring service as long as it delivers the measurement information in NMWG XML format.

### **5.2.3 Measurement archive service**

A measurement archive stores measurement information in a database. This measurement data that is produced by a measurement point is also archived. Historical data can be requested from a measurement archive.

The measurement archive service must be subscribed to a measurement point service or another measurement archive service to get measurement data. It also must be registered with the lookup service, so the client is able to find the archive. A client is able to access the measurement data directly after requesting the lookup service for its location.

### **5.2.4 Authentication service**

The authentication service is not yet implemented in PerfSONAR. This service should control the authentication of the users and/or group. Whether they have access to measurement data, or if the user can control a measurement point, etc.

This service communicates with the resource protector service (subsection 5.2.6).

### **5.2.5 Transformation service**

The transformation service performs a function upon the measurement data. This can be an aggregation, correlation, translation, etc. An example of this service is the topology service, which provides necessary information about the network. There is not yet a transformation service implemented in the PerfSONAR framework.

A topology service collects information from a variety of sources, e.g. measurement points and measurement archives. Algorithms derive the network topology from this information.

### **5.2.6 Resource protector service**

The Resource Protector Service manages the authorization of users and the resource availability requests. There is not yet an implementation of this service available, it is only documented.

## **5.3 Communication between services**

PerfSONAR is service oriented architecture, therefore it communication between the services with Simple Object Access Protocol (SOAP). This are XML messages over the HTTP protocol. This communication is done according the NMWG schema. This is a XML schema developed for network measurements.

In the NMWG schema the topology points, the links and link status are identified. Examples are given in appendix A. In this report the NMWG schema will not be explained in detail, this schema is described in detail in [8] and [17].

## 6 E2E inter-domain monitoring in the GLIF infrastructure

Currently monitoring E2E (end-to-end) links within a single domain is not a problem, but as described in section 2 the challenge lies in the inter-domain aspect of monitoring an E2E link. To recapitulate: In figure 1 each NOC monitors the lightpaths within its own domain. However, E2E lightpaths often cross several domains and especially in the GLIF infrastructure. If an E2E lightpath breaks, the NOCs can only see if the problem resides in their own domain. To verify where the problem exists the NOCs need to contact each other which can be very time consuming. The next subsections cover the E2E Link Monitoring System, developed by GN2-JRA<sup>1,2</sup>, which solves the inter-domain monitoring problem by providing the intra-domain monitoring data to the other domains.

### 6.1 E2E Monitoring System

The E2E Link Monitoring System, developed by GEANT2, is in essence a visualization tool written in HTML. It is able to visualize the status of an inter-domain E2E link as depicted in figure 6. However in order to visualize the status of an E2E inter-domain link it is necessary to get all the statuses of the separate links the E2E link is composed of. The E2E Monitoring System explains the concepts of E2E Monitoring. It states what each domain has to do in order to participate (e.g. how measurement data should be provided to the E2E Monitoring System).

**Status of E2E Link CERN-SARA-LHCOPN-001**

Time of State Aggregation: 2007-06-28, 14:08:18 MEST (Cycle time: 60 s.)  
Operational State: **Up**  
Administrative State: **Normal Oper.**

Domain	GEANT2					CERN (?)			
Link Structure	EP	←-----	-----→	DP	←-----→	DP	←-----	-----→	EP
Type	EndPoint	ID Part.Info	ID Part.Info	Demarc	Domain Link	Demarc	ID Part.Info	ID Part.Info	EndPoint
Local Name	SURFNET-SARA	SURFNET-NL-CERN-SARA	SURFNET-LHC_ter1.15.ams-fra.ams	GEANT2-AMS	ams-gen_LHC_CERN-SURFNET_06002	GEANT2-GEN	CERN-LHC_ext1.18.fra-gen.gen	S513-C-BE3	CERN-TO
State Oper.	-	<b>Up</b>	<b>Up</b>	-	<b>Up</b>	-	<b>Up</b>	<b>Up</b>	-
State Admin.	-	<b>Normal Oper.</b>	<b>Normal Oper.</b>	-	<b>Normal Oper.</b>	-	<b>Normal Oper.</b>	<b>Normal Oper.</b>	-
Timestamp	-	2007-06-28 T14:05:08.0+0200	2007-06-28 T13:07:20.0+0100	-	2007-06-28 T13:07:20.0+0100	-	2007-06-28 T13:07:20.0+0100	2007-06-28 T14:07:53+02:00	-

Figure 6: Visualization of an E2E link

### 6.2 Concepts

In [5] an E2E link is defined as an optical connection realized using layer 1 and 2 technologies like Ethernet or SDH/SONET over fiber. Such a link connects two endpoints that may reside in different domains (see figure 7). The E2E link is composed of several local links and inter-domain links. A local link is a link within a single domain. An inter-domain link connects two domains. These two link types are called “logical links” for the remainder of this document. A node that is not the end of an E2E link is called a demarcation point. The

<sup>2</sup>GEANT2 Joint Research Activity 1

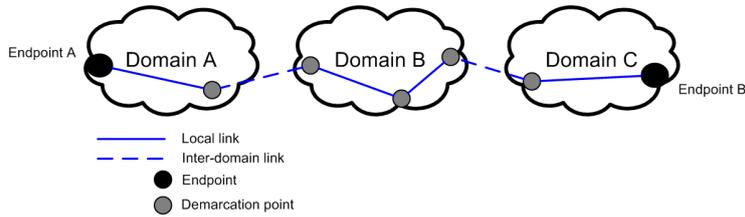


Figure 7: Concept of an E2E inter-domain link

node ending an E2E link is obviously called an endpoint. Thus, an E2E link is a logical connection between two endpoints composed of one or more logical links.

The technology used on these links (e.g. SDH/SONET, Ethernet, etc.) is irrelevant for the E2E Monitoring System as they are treated as logical links which share common abstract properties (e.g. link status). The E2E Monitoring System polls the domains' Measurement Archives to retrieve the status information on all the links (see figure 8). It only needs to access the MA Web Service and not the physical devices. In order to compose a correct E2E link a naming convention is used by all involved parties. An E2E link is owned by a

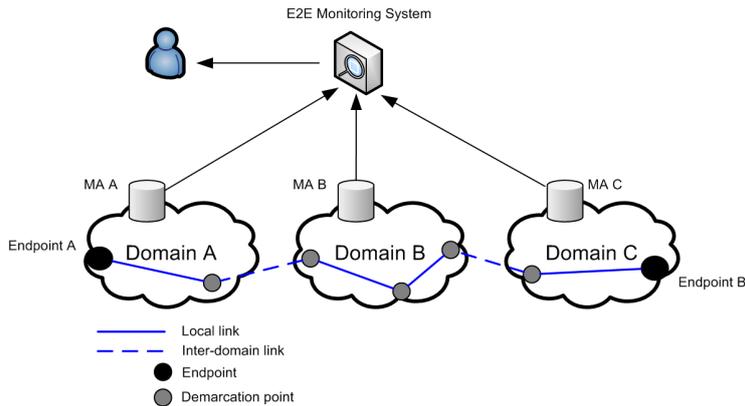


Figure 8: E2E Link Monitoring System

project and has two state attributes representing the operational and the administrative status of the complete E2E link. The possible values of the attributes are respectively presented in table 1 and 2. A logical link has three attributes, including the two previously mentioned attributes. The third attribute is a time stamp which indicates the last state measurement of the link. A Topology Point is either an endpoint or a demarcation point and it is part of an authoritative domain. This domain is responsible for the Topology Point. A link consists of at least two Topology Points. A Topology Point has the following attributes: Name, Country, City, Institution, Latitude and Longitude. To retrieve the status of an inter-domain link two scenarios are possible. In the first scenario one domain takes the responsibility for monitoring the link. In the second scenario, which is in practice also the most likely to occur, both domains monitor their side of the link. Information of both domains is then aggregated. This type of

Operational State Value	Description	Weight
Unknown	Domain could not acquire information about operational state	0
Up	Link is up	1
Degraded	Link is up, but has degraded performance	2
Down	Link is down	3

Table 1: Operational state [5]

Operational State Value	Description	Weight
Unknown	Domain could not acquire information about operational state	0
Normal operation	No administrative work is performed	1
Maintenance	Planned maintenance activity in progress	2
Troubleshooting	Trouble shooting is in progress	3
UnderRepair	Repair process is in progress	4

Table 2: Administrative state [5]

link is called “Inter Domain Partial Info”.

### 6.2.1 Example of an E2E inter-domain link

The configuration files in appendix A are an example of how an E2E inter-domain link in PerfSONAR is defined. The global link name that is used for this lightpath is: CERN-SARA-LHCOPN-001. The global link name consists

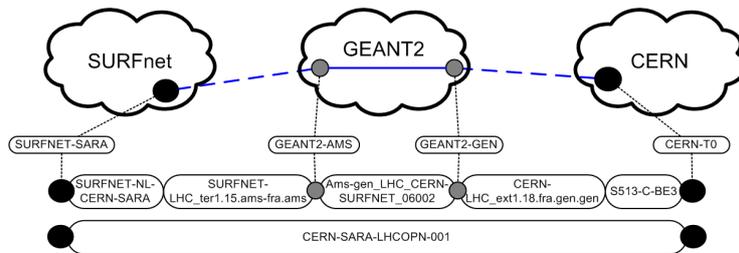


Figure 9: Example of an E2E inter-domain link

of four parts. The first two parts are the domain names which are connected via the lightpath. The third part is the name of the project the link belongs to and the fourth part is a sequence number. If multiple lightpaths are used for the same project between the same nodes, this number distinguishes the different lightpaths.

### 6.2.2 E2E link reconstruction and aggregating link information

Before the E2E Monitoring System can provide the status of an E2E link it has to reconstruct the link by aggregating the logical links in the correct order.

This process is possible only when two endpoints are given and the naming of endpoints and demarcation points is done consistently across the involved domains.

When providing status information about an E2E link the status is equal to the status of the link with the highest weight (see 1 and 2).

### 6.3 Data provisioning

Figure 10 shows a scenario for the data provisioning to feed the E2E Monitoring System. First a network monitoring system retrieves the status information from the equipment. NetherLight<sup>3</sup> and SURFnet6<sup>4</sup> currently use the TL1 toolkit to do this but a new tool based on the TL1 toolkit, called Spotlight (see 6.4) is developed.

In order to deliver data to the MA web service correctly the raw data from the devices has to be parsed in a way it complies with the data model.

When the data is stored in the database it is accessible through the MA. The E2E Monitoring System is now able to retrieve the needed information and display the status of E2E links.

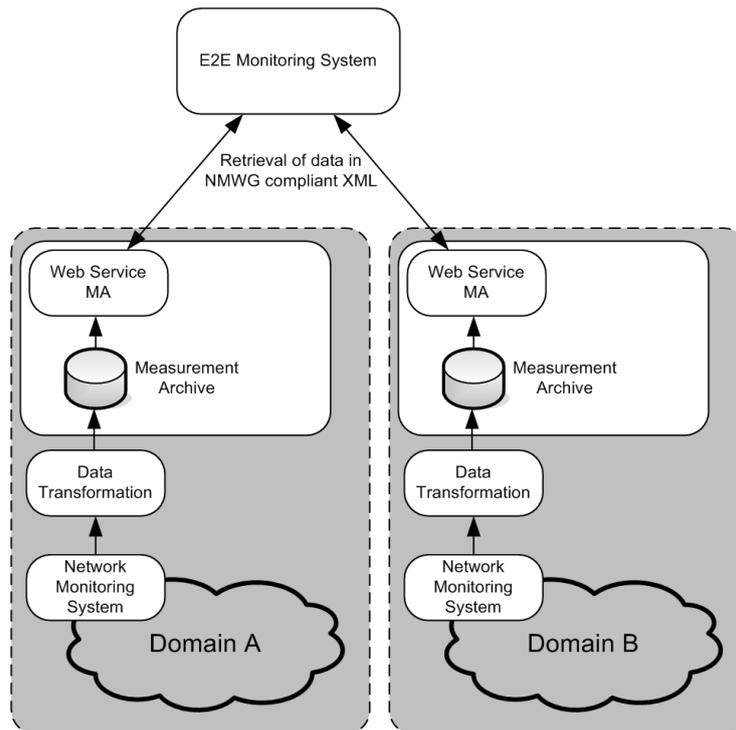


Figure 10: Data provisioning with the use of the Measurement Archive Web Service

<sup>3</sup>NetherLight is an optical exchange of the GLIF infrastructure.

<sup>4</sup>SURFnet6 is the National Research and Education Network of the Netherlands.

## 6.4 PerfSONAR and Spotlight

Currently the monitoring of the SURFnet6 and NetherLight architecture is done by the TL1 toolkit which is part of Spotlight. It retrieves network information (e.g. capacity, free timeslots and alarm types) from the hardware using TL1 commands and stores it in a database. This toolkit also generates a Network Description Language (NDL)[18] compliant file to describe the network topology. For monitoring E2E links such a topology description is required in order to know how the E2E link is build up.

The PerfSONAR framework describes the topology with a NMWG compliant file. Furthermore it defines which data is measured and where it can be found in the database. If NetherLight wants to participate and provide network information to other domains with the use of PerfSONAR, the already deployed TL1 toolkit can be used to 1) acquire information, 2) fill the network status database with this information and 3) generate a NMWG compliant file (see figure 11). The MA web service should be modified to retrieve data from the

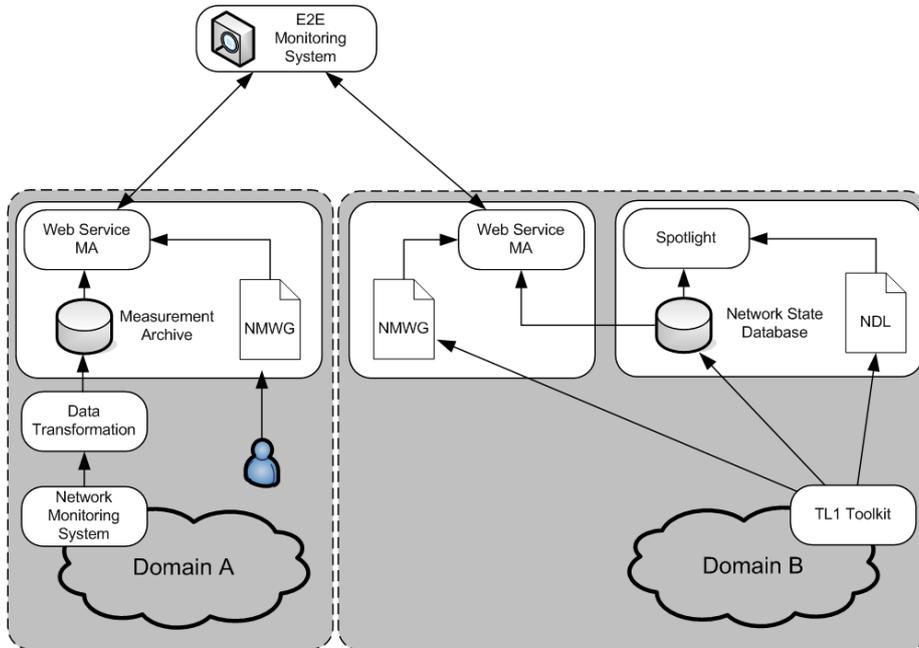


Figure 11: Proposed inter relationship between Spotlight and PerfSONAR

network status database. The reason for using the MA web service and not modifying Spotlight in a way it can replace the MA web service is that modifications and added functionality in the MA webservice will also need to be implemented in Spotlight. This makes it hard and time consuming to maintain. The advantage of integrating Spotlight with PerfSONAR is the fact that the information retrieved from the network by the TL1 toolkit is made available through the web services of the PerfSONAR framework which is already used by multiple organizations.

## 7 Future work

This report explains the possible use of PerfSONAR in the GLIF infrastructure for monitoring E2E links. However this research raised some other question.

One of the requirements for the GLIF infrastructure is the support for monitoring protected lightpaths. Currently this is not possible. Future work could include doing research on how to support protected lightpaths in PerfSONAR.

PerfSONAR is focussed on performance monitoring. Instead of performance monitoring, this study is focussed on E2E lightpath monitoring. However a study on performance monitoring with the use of PerfSONAR can be valuable for the StarPlane network.

Lightpaths in the GLIF infrastructure change on a regular basis. In PerfSONAR static NMWG files are used which are needed to describe the network topology and which data is being measured. This file is currently manually created. Tools that generate that file when the topology changes are needed to make maintenance less time consuming.

## 8 Conclusions

This project's primary objective was to research the possible use of the PerfSONAR framework for inter-domain monitoring in the GLIF infrastructure. The different stakeholders were identified and their requirements defined (see 4.2). PerfSONAR is currently capable of monitoring inter-domain E2E lightpaths. However, there are unmet requirements that make PerfSONAR less suitable for the GLIF infrastructure. Disadvantages are; the lack of support for protected and redundant lightpaths. Authentication and authorization within the framework is not yet implemented. The number of supported metrics is very limited. Advantages are; the architecture of PerfSONAR allows for additional metrics. The decentralized character of PerfSONAR improves scalability. Extending the PerfSONAR framework with customized services is possible due to its service oriented architecture which makes it flexible.

## Acknowledgments

We would like to thank Ronald van der Pol, Andree Toonk and Paola Grosso for their help and effort they put into this research project, their valuable comments on previous versions of this paper and comments on the presentation. Also we would like to thank the people in the PerfSONAR mailing list for their comments and their fast replies. Last but not least we thank SARA who gave us the opportunity to do this project.

## References

- [1] Jeff W. Boote, Eric L. Boyd, Jérôme Durand, Andreas Hanemann, Loukik Kudarimoti, Roman Lapacz, Nicolas Simar, and Szymon Trocha. Towards multi-domain monitoring for the european research networks. *Computational methods in science and technology*, 11(2):91–100, 2005.
- [2] ESnet, Energy Sciences network. <http://www.es.net>, April 2007.
- [3] GEANT2, Gigabit European Academic Network 2. <http://www.geant2.net>, June 2007.
- [4] GLIF, Global Lambda Integrated Facility. <http://www.glif.is>, May 2007.
- [5] Matthias Hamm and Mark Yampolskiy. E2e link monitoring: System design and documentation. page 57, 2007.
- [6] A. Hanemann, V. Jeliaskov, O. Kvittem, L. Marta, J. Metzger, and I. Velimirovic. Complementary visualization of perfSONAR network performance measurements. *ICISP*, 0:6, 2006.
- [7] Andreas Hanemann, Jeff W. Boote, Eric L. Boyd, Jérôme Duran, Loukik Kudarimoti, Roman Lapacz, D. Martin Swany, Szymon Trocha, and Jason Zurawski. PerfSONAR : A service oriented architecture for multi-domain network monitoring. *LNCS*, 3826:241–254, 2005.
- [8] Andreas Hanemann, Loukik Kudarimoti, Roman Lapacz, D. Martin Swany, Szymon Trocha, and Jason Zurawski. Application of the NM-WG schema for network performance data exchange in the perfSONAR system.
- [9] René Hatem, Almar Giesberts, and Erik-Jan Bos. The ordering and fault resolution process for multi-domain lightpaths across hybrid networks. 0:5, 2006.
- [10] Internet2. <http://www.internet2.edu>, June 2007.
- [11] Internet2. Bandwidth test controller (BWCTL). <http://e2epi.internet2.edu/bwctl/>, September 2005.
- [12] University of Illinois. <http://dast.nlanr.net/Projects/Iperf/>, 2005.
- [13] Mike P. Papazoglou. Service -oriented computing: Concepts, characteristics and directions. In *WISE '03: Proceedings of the Fourth International Conference on Web Information Systems Engineering*, page 3, Washington, DC, USA, 2003. IEEE Computer Society.
- [14] perfSONAR, PERFORMANCE focused Service Oriented Network monitoring ARchitecture. <http://www.perfsonar.net>, June 2007.
- [15] RNP, Rede Nacional de ensino e Pesquisa. <http://www.rnp.br/en>, May 2007.
- [16] SURFnet. Measuring the use of lightpath services. 0:13, 2007.

- [17] D. Martin Swany, Jason W. Zurawski, and Dan Gunter. NMWG schema developers guide. August, 2005.
- [18] Research Group System and Network Engineering. Network description language. <http://www.science.uva.nl/research/sne/ndl/>, 2006.
- [19] XQuery, an XML query language. <http://www.w3.org/TR/xquery>, January 2007.

## A XML E2E-link configuration files

These configuration files are an example of an E2E inter-domain link. The global link name that is used for this lightpath is: CERN-SARA-LHCOPN-001.

### A.1 SURFnet domain

```
<?xml version="1.0" encoding="utf-8" ?>
<nmwg:store type="E2E_Link_status_information"
  xmlns:nmwg="http://ggf.org/ns/nmwg/base/2.0/"
  xmlns:nmtm="http://ggf.org/ns/nmwg/time/2.0/"
  xmlns:nmwgtopo3="http://ggf.org/ns/nmwg/topology/base/3.0/"
  xmlns:nmtl2="http://ggf.org/ns/nmwg/topology/l2/3.0/"
  xmlns:nmtl3="http://ggf.org/ns/nmwg/topology/l3/3.0/"
  xmlns:nmtl4="http://ggf.org/ns/nmwg/topology/l4/3.0/"
  xmlns:ifevt="http://ggf.org/ns/nmwg/event/status/base/2.0/">

  <nmwg:parameters id="storeId">
    <nmwg:parameter name="DomainName">SURFnet</nmwg:parameter>
  </nmwg:parameters>

  <nmwg:metadata id="md1">
    <nmwg:subject id="sub-SURFNET-SARA">
      <nmwgtopo3:node id="SURFNET-SARA">
        <nmwgtopo3:type>TopologyPoint</nmwgtopo3:type>
        <nmwgtopo3:name type="logical">
          SURFnet-Asd001A
        </nmwgtopo3:name>
        <nmwgtopo3:country>
          NL
        </nmwgtopo3:country>
        <nmwgtopo3:city>Amsterdam</nmwgtopo3:city>
        <nmwgtopo3:institution>
          Asd001A
        </nmwgtopo3:institution>
        <nmwgtopo3:latitude>
          node1_latitude
        </nmwgtopo3:latitude>
        <nmwgtopo3:longitude>
          node1_longitude
        </nmwgtopo3:longitude>
      </nmwgtopo3:node>
    </nmwg:subject>
  </nmwg:metadata>

  <nmwg:metadata id="md2">
    <nmwg:subject id="sub1">
      <nmtl2:link>
        <nmtl2:name type="logical">
          SURFNET-NL-CERN-SARA
        </nmtl2:name>
      </nmtl2:link>
    </nmwg:subject>
  </nmwg:metadata>
</nmwg:store>
```

```

        </nmtl2:name>
        <nmtl2:globalName type="logical">
            CERN-SARA-LHCOPN-001
        </nmtl2:globalName>
        <nmtl2:type>ID_LinkPartialInfo</nmtl2:type>
        <nmwgtopo3:node nodeIdRef="SURFNET-SARA">
            <nmwgtopo3:role>
                EndPoint
            </nmwgtopo3:role>
        </nmwgtopo3:node>
        <nmwgtopo3:node nodeIdRef="GEANT2-AMS">
            <nmwgtopo3:role>
                DemarcPoint
            </nmwgtopo3:role>
        </nmwgtopo3:node>
    </nmtl2:link>
</nmwg:subject>

    <nmwg:parameters>
        <nmwg:parameter name="supportedEventType">
            UP
        </nmwg:parameter>
    </nmwg:parameters>
</nmwg:metadata>

    <nmwg:data id="d1" metadataIdRef="md2">
        <ifevt:datum timeType="ISO" timeValue="2007-06-29T00:05:07.0+0200">
            <ifevt:stateOper>UP</ifevt:stateOper>
            <ifevt:stateAdmin>NORMALOPERATION</ifevt:stateAdmin>
        </ifevt:datum>
    </nmwg:data>

</nmwg:store>

```

## A.2 GEANT2 domain

```

<?xml version="1.0" encoding="utf-8"?>
<nmwg:message type="E2E_Link_status_information"
    xmlns:nmwg="http://ggf.org/ns/nmwg/base/2.0/"
    xmlns:nmtm="http://ggf.org/ns/nmwg/time/2.0/"
    xmlns:nmwgtopo3="http://ggf.org/ns/nmwg/topology/base/3.0/"
    xmlns:nmtl2="http://ggf.org/ns/nmwg/topology/l2/3.0/"
    xmlns:nmtl3="http://ggf.org/ns/nmwg/topology/l3/3.0/"
    xmlns:nmtl4="http://ggf.org/ns/nmwg/topology/l4/3.0/"
    xmlns:ifevt="http://ggf.org/ns/nmwg/event/status/base/2.0/">
    <nmwg:parameters id="storeId">
        <nmwg:parameter name="DomainName">GEANT2</nmwg:parameter>
    </nmwg:parameters>

```

```

<nmgw:metadata id="md1">
  <nmgw:subject id="sub-GEANT2-GEN">
    <nmgwtopo3:node id="GEANT2-GEN">
      <nmgwtopo3:type>TopologyPoint</nmgwtopo3:type>
      <nmgwtopo3:name type="logical">
        GEANT2-GEN
      </nmgwtopo3:name>
      <nmgwtopo3:country>
        CH
      </nmgwtopo3:country>
      <nmgwtopo3:city>Geneva</nmgwtopo3:city>
      <nmgwtopo3:institution>
        Cern
      </nmgwtopo3:institution>
      <nmgwtopo3:latitude>
        node1_latitude
      </nmgwtopo3:latitude>
      <nmgwtopo3:longitude>
        node1_longitude
      </nmgwtopo3:longitude>
    </nmgwtopo3:node>
  </nmgw:subject>
</nmgw:metadata>

<nmgw:metadata id="md2">
  <nmgw:subject id="sub-GEANT2-AMS">
    <nmgwtopo3:node id="GEANT2-AMS">
      <nmgwtopo3:type>TopologyPoint</nmgwtopo3:type>
      <nmgwtopo3:name type="logical">
        GEANT2-AMS
      </nmgwtopo3:name>
      <nmgwtopo3:country>
        NL
      </nmgwtopo3:country>
      <nmgwtopo3:city>Amsterdam</nmgwtopo3:city>
      <nmgwtopo3:institution>
        Sara
      </nmgwtopo3:institution>
      <nmgwtopo3:latitude>
        node2_latitude
      </nmgwtopo3:latitude>
      <nmgwtopo3:longitude>
        node2_longitude
      </nmgwtopo3:longitude>
    </nmgwtopo3:node>
  </nmgw:subject>
</nmgw:metadata>

<nmgw:metadata id="md3">
  <nmgw:subject id="sub1">

```

```

<nmtl2:link>
  <nmtl2:name type="logical">
    ams-gen_LHC_CERN-SURFNET_06002
  </nmtl2:name>
  <nmtl2:globalName type="logical">
    CERN-SARA-LHCOPN-001
  </nmtl2:globalName>
  <nmtl2:type>DOMAIN_Link</nmtl2:type>
  <nmwgtopo3:node nodeIdRef="GEANT2-GEN">
    <nmwgtopo3:role>
      DemarcPoint
    </nmwgtopo3:role>
  </nmwgtopo3:node>
  <nmwgtopo3:node nodeIdRef="GEANT2-AMS">
    <nmwgtopo3:role>
      DemarcPoint
    </nmwgtopo3:role>
  </nmwgtopo3:node>
</nmtl2:link>
</nmwg:subject>
<nmwg:parameters>
  <nmwg:parameter name="supportedEventType">
    UP
  </nmwg:parameter>
</nmwg:parameters>
</nmwg:metadata>

<nmwg:metadata id="md4">
  <nmwg:subject id="sub2">
    <nmtl2:link>
      <nmtl2:name type="logical">
        CERN-LHC_ext1.18.fra-gen.gen
      </nmtl2:name>
      <nmtl2:globalName type="logical">
        CERN-SARA-LHCOPN-001
      </nmtl2:globalName>
      <nmtl2:type>ID_LinkPartialInfo</nmtl2:type>
      <nmwgtopo3:node nodeIdRef="GEANT2-GEN">
        <nmwgtopo3:role>
          DemarcPoint
        </nmwgtopo3:role>
      </nmwgtopo3:node>
      <nmwgtopo3:node nodeIdRef="CERN-T0">
        <nmwgtopo3:role>
          EndPoint
        </nmwgtopo3:role>
      </nmwgtopo3:node>
    </nmtl2:link>
  </nmwg:subject>
<nmwg:parameters>

```

```

        <nmgw:parameter name="supportedEventType">
            UP
        </nmgw:parameter>
    </nmgw:parameters>
</nmgw:metadata>

<nmgw:metadata id="md5">
    <nmgw:subject id="sub3">
        <nmtl2:link>
            <nmtl2:name type="logical">
                SURFNET-LHC_ter1.15.ams-fra.ams
            </nmtl2:name>
            <nmtl2:globalName type="logical">
                CERN-SARA-LHCOPN-001
            </nmtl2:globalName>
            <nmtl2:type>ID_LinkPartialInfo</nmtl2:type>
            <nmgwtopo3:node nodeIdRef="GEANT2-AMS">
                <nmgwtopo3:role>
                    DemarcPoint
                </nmgwtopo3:role>
            </nmgwtopo3:node>
            <nmgwtopo3:node nodeIdRef="SURFNET-SARA">
                <nmgwtopo3:role>
                    EndPoint
                </nmgwtopo3:role>
            </nmgwtopo3:node>
        </nmtl2:link>
    </nmgw:subject>
    <nmgw:parameters>
        <nmgw:parameter name="supportedEventType">
            UP
        </nmgw:parameter>
    </nmgw:parameters>
</nmgw:metadata>

<nmgw:data id="d1" metadataIdRef="md3">
    <ifevt:datum timeType="ISO" timeValue="2007-06-27T15:39:46.0+0100">
        <ifevt:stateOper>UP</ifevt:stateOper>
        <ifevt:stateAdmin>NORMALOPERATION</ifevt:stateAdmin>
    </ifevt:datum>
</nmgw:data>

<nmgw:data id="d2" metadataIdRef="md4">
    <ifevt:datum timeType="ISO" timeValue="2007-06-27T15:39:46.0+0100">
        <ifevt:stateOper>UP</ifevt:stateOper>
        <ifevt:stateAdmin>NORMALOPERATION</ifevt:stateAdmin>
    </ifevt:datum>
</nmgw:data>

<nmgw:data id="d3" metadataIdRef="md5">

```

```

        <ifevt:datum timeType="ISO" timeValue="2007-06-27T15:39:46.0+0100">
            <ifevt:stateOper>UP</ifevt:stateOper>
            <ifevt:stateAdmin>NORMALOPERATION</ifevt:stateAdmin>
        </ifevt:datum>
    </nmwg:data>

</nmwg:message>

```

### A.3 CERN domain

```

<?xml version="1.0" encoding="utf-8" ?>
<nmwg:message type="response"
  xmlns:nmwg="http://ggf.org/ns/nmwg/base/2.0/"
  xmlns:nmwgr="http://ggf.org/ns/nmwg/result/2.0/"
  xmlns:nmwgt="http://ggf.org/ns/nmwg/topology/2.0/"
  xmlns:nmwgtopo3="http://ggf.org/ns/nmwg/topology/base/3.0/"
  xmlns:nmtl2="http://ggf.org/ns/nmwg/topology/l2/3.0/"
  xmlns:nmtm="http://ggf.org/ns/nmwg/time/2.0/"
  xmlns:ifevt="http://ggf.org/ns/nmwg/event/status/base/2.0/"
  xmlns:netutil="http://ggf.org/ns/nmwg/characteristics/utilization/2.0/">

  <nmwg:metadata id="S513-C-BE3">
    <nmwg:subject id="S513-C-BE3">
      <nmtl2:link>
        <nmtl2:name type="logical">
          S513-C-BE3
        </nmtl2:name>
        <nmtl2:globalName type="logical">
          CERN-SARA-LHCOPN-001
        </nmtl2:globalName>
        <nmtl2:type>ID_LinkPartialInfo</nmtl2:type>
        <nmwgtopo3:node nodeIdRef="CERN-T0">
          <nmwgtopo3:role>
            EndPoint
          </nmwgtopo3:role>
        </nmwgtopo3:node>
        <nmwgtopo3:node nodeIdRef="GEANT2-GEN">
          <nmwgtopo3:role>
            DemarcPoint
          </nmwgtopo3:role>
        </nmwgtopo3:node>
      </nmtl2:link>
    </nmwg:subject>
  </nmwg:metadata>

  <!--the node with ID "CERN-T0" is missing!
        it should look something like this: -->

  <nmwg:metadata id="md1">
    <nmwg:subject id="sub-CERN-T0">

```

```

    <nmwgtopo3:node id="CERN-T0">
      <nmwgtopo3:type>TopologyPoint</nmwgtopo3:type>
      <nmwgtopo3:name type="logical">
        CERN-T0
      </nmwgtopo3:name>
      <nmwgtopo3:country>
        CH
      </nmwgtopo3:country>
      <nmwgtopo3:city>Geneva</nmwgtopo3:city>
      <nmwgtopo3:institution>
        Cern
      </nmwgtopo3:institution>
      <nmwgtopo3:latitude>
        node1_latitude
      </nmwgtopo3:latitude>
      <nmwgtopo3:longitude>
        node1_longitude
      </nmwgtopo3:longitude>
    </nmwgtopo3:node>
  </nmwg:subject>
</nmwg:metadata>

<!--End example -->

<nmwg:data metadataIdRef="S513-C-BE3" id="4335214">
  <ifevt:datum timeType="ISO" timeValue="2007-06-27T16:40:39+02:00">
    <ifevt:stateOper>Up</ifevt:stateOper>
    <ifevt:stateAdmin>NormalOperation</ifevt:stateAdmin>
  </ifevt:datum>
</nmwg:data>

</nmwg:message>

```