

Global Lightpath Identifiers Proposal

Lars Fischer Tom Lehman Ronald van der Pol Thomas Tam

Version 1.2
August, 2008

1 Introduction

Currently, each domain uses its own identifiers for lightpaths that span multiple domains. This makes it difficult to refer to the same lightpath during the provisioning phase, in case of outages or when announcing planned work.

At the GLIF Working Group Meeting in January 2008 the GLIF community decided to set up a task force to work on a global lightpath identifier scheme. The task force consists of Ronald van der Pol (leader), Lars Fischer, Tom Lehman and Thomas Tam.

Global identifiers complement the local naming schemes that are in use in the various domains. It is assumed that most domains will use the global identifiers as aliases for their local names. The global identifiers are used in communication with other domains.

In this document a couple of proposals are described. We have defined a couple of requirements that can be used to compare the various naming schemes. These requirements are described in section 2. The recommendation of the task force is described in section 4.

2 Requirements

The following requirements are used to compare the various naming schemes in section 3.

- The global identifier should preferably consist of a unique identifier for a certain domain followed by a local part that is unique within that domain. Each domain can choose its own naming scheme for the local part. This was the preference of many people during the GLIF meeting in January 2008.
- Centralized registries should be avoided. The GLIF community uses a distributed approach to operations. Each domain has full control over its own operations. Centralized registries raise difficult questions about where to put them, who runs them, who pays for them, etc.
- A maximum length of the identifier should be set. Very long names are inconvenient. Too long names limit the places where the identifiers can be used, stored and/or displayed.
- The allowed character set for global identifiers should be limited to a well-defined character set. This makes them easy to parse with regular expressions. It also avoids surprises when somebody wants to use a global identifier in a system that has a special meaning for certain special characters.

3 Naming Schemes Examples

Section 3.1 and 3.2 are naming schemes that are already in use in real life. Sections 3.3 to 3.5 are proposals that are not being used yet.

3.1 DANTE Naming Scheme

In projects like the LHCOPN and DEISA a naming scheme for the end-to-end monitoring of lightpaths with perfSONAR is used. It consists of the two end points, the project name and a sequence number. E.g., CERN-TRIUMF-LHCOPN-001. In this naming scheme a couple of names have to be agreed upon. Each end site must have a globally unique identifier and a central registry with these identifiers must be set up. The same is true for the project names. The sequence number also needs a central registry in order to find the next available sequence number.

3.2 Internet2 Naming Scheme

This scheme provides for a unique global identifier for each instance of a circuit that is reserved or provisioned. This identifier is referred to as a "Global Reservation Identifier" (GRI). This GRI is established at the time of reservation and utilized throughout the lifetime of the circuit. It can also be used by other processes, such as monitoring systems, in order to query for additional information.

The process for establishment of a GRI is as follows:

- The domain which initiates a service request is responsible for establishment of the globally unique GRI.
- The GRI is constructed by combining two parts, *domainid*-*serviceinstanceid*, where:
 - *domainid*: DNS style name which utilizes the initiating domain DNS name structure.
 - *serviceinstanceid*: A number generated by the service initiating domain, which is unique within the scope of that domain. This could be a statistically unique number so that the domain does not need to setup a registry for *serviceinstanceid* numbers.

The above scheme allows for the initiating domain to create a globally unique identification. Below are examples of GRIs generated by Internet2 and ESnet which are in accordance with this scheme.

- dcn.internet2.edu-6811
- es.net-4005

Where the numerical part is the *serviceinstanceid*, and is managed by the individual domains to ensure uniqueness within their domain.

This naming scheme is currently used by Internet2 and was jointly developed by several organizations, including Internet2, ESNNet, GEANT, and others.

3.3 URN Naming Scheme

Uniform Resource Names (URNs) are defined in RFC2141 where they are described as *persistent, location-independent, resource identifiers*. The format of URNs is:

```
<URN> ::= "urn:" <NID> ":" <NSS>
```

where <NID> is the Namespace Identifier, and <NSS> is the Namespace Specific String.

For GLIF global identifiers the NID is set to *glif*. The <NSS> is a variant of the Internet2 naming scheme. The first part of the <NSS> is the *sourcing organisation* identifier (instead of a DNS domain name as in the Internet2 Naming Scheme). The sourcing organisation is described in [2]. The second part is set by the sourcing organisation and is unique within that domain. An example of this naming scheme is *urn:glif:starlight:chi-van-42*.

3.4 Other Hierarchical Naming Schemes

This needs a centralised repository for GOLE abbreviations. This list could be kept on the GLIF wiki, which already has a list of all the GOLEs.

Other possibilities for the domain part are city abbreviations, like IATA airport codes or UN LOCODEs [1]. The city abbreviation used is the city where the institution of one of the end points of the lightpath is situated. A disadvantage of IATA codes is that these only work for cities with an airport. This problem does not exist for UN LOCODEs, but finding the correct LOCODE can also be difficult in some cases. This is especially true for small villages and these are typically the places where radiotelescopes are built.

3.5 Flat Naming Schemes

Another type of identifiers are strings without hierarchy. These could be unique numbers or unique alpha-numeric strings. These identifiers can be kept in a central registry to make sure that they are unique.

Another possibility is to have statistically unique identifiers. In this case the chance of randomly picking an identifier that is already in use is very small. When using a 8 character alpha-numeric string, the chances of a clash are 1 in 2,821,109,907,456. When there are 500,000 lightpaths in use in GLIF, the chance of picking a new identifier that is already in use is 1 in 5,642,220.

4 Recommendation

We propose to use the *URN Naming Scheme* for global identifiers. This is a standardised format (RFC2141) and fits well in web services frameworks. The identifiers are set in a distributed way in which the sourcing organisation of a lightpath creates the identifier. This follows the lightpath setup and fault resolution procedures within the GLIF community, as described in [2], most closely. Therefore, we expect that this naming scheme is easy to implement by the GLIF community. A *sourcing organisation* can be an NREN, Regional Network, GOLE, etc.

Global identifiers are defined as follows:

```
<global identifier> ::= "urn:glif:" <NSS>
```

The *sourcing organisation* generates the <NSS> by using the following rules:

- The sourcing organisation picks a globally unique identifier. Using its DNS name is an example of how this can be accomplished. This identifier is the first part of the <NSS>.

- The second part of the <NSS> is set to some internally unique naming scheme.
- The first and second part are separated by a colon ":". There should be only one ":" in the entire <NSS>.
- The format of the first and second part of the <NSS> are determined by the sourcing organisation. There is only a restriction with respect to the allowed character set as follows:
 - only alpha-numeric, hyphen (-) and period (.) are allowed
 - no blank or space characters are allowed
 - no distinction is made between upper and lower case

The total length of the global identifier should be at most 64 characters.

These rules satisfy the requirements in section 2. The URN based global identifiers contain a domain identifier followed by a local part. There are no centralised registries needed. These global identifiers also have a fixed length of at most 64 characters and the allowed character set is well-defined.

Examples are:

- urn:glif:dcn.internet2.edu:5705
- urn:glif:dragonmax:5702
- urn:glif:starlight:chi-van-42
- urn:glif:northernlight:G6O76GQ
- urn:glif:cern:LHCOPN-CA-TRIUMF-CH-CERN-002

References

- [1] <http://en.wikipedia.org/wiki/UN/LOCODE>.
- [2] Rene Hatem, Almar Giesberts, and Erik-Jan Bos, *The ordering and fault resolution process for multi-domain lightpaths across hybrid networks*, July 2006, <http://www.glif.is/working-groups/tech/fault-resolution-0.9.pdf>.