

# Global Lightpath Identifiers Proposal

Lars Fischer      Tom Lehman      Ronald van der Pol      Thomas Tam

Draft 0.1  
April, 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 criteria that can be used to compare the various naming schemes. These criteria are described in section 3. The recommendation of the task force is described in section 4.

## 2 Examples Naming Schemes

### 2.1 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.

### 2.2 Internet2 Naming Scheme

The naming scheme used in Internet2 consists of two parts. The identifier starts with the DNS domain name of one of the end points. (Tom, can you correct and expand this part? What does the domain local part look like?)

## 2.3 Hierarchical Naming Schemes

During the GLIF meeting in January many working group members were in favour of a scheme that is similar to the one used by Internet2. So, a name consisting of a unique domain identifier, followed by a second part that is unique within that domain.

A unique domain identifier could be an abbreviation for a GOLE. This will be the originating GOLE for that lightpath. These GOLE abbreviations will be kept in a central registry, e.g. the GLIF website. Each GOLE uses a naming scheme for the second part that uniquely identifies each lightpath.

Other possibilities for the domain part are city abbreviations, like IATA airport codes or UN LOCODEs [1].

## 2.4 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.

# 3 Criteria

The following criteria are used to compare the various naming schemes.

1. Is a global registry needed?
2. Is a central registry per domain needed?
3. Does the identifier have a fixed length?
4. Can the identifier be generated by provisioning software?
5. Is the identifier unique or only statistically unique?

Other things to consider are the allowed characters in the name. The global identifiers may end up in various places, like management systems, interface descriptions, databases, (dynamic) provisioning systems, path or section traces, VLAN names, etc. Therefore, it seems wise to be conservative and restrict the allowed character set to the rules for ARPANET host names. They must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and hyphen.

It also seems wise to restrict the names to a certain maximum length. Otherwise they will be cut off when using them for one of the things mentioned in the previous paragraph. After cut off the identifier may not be unique anymore.

## 4 Recommendation

Global or domain registries have several drawbacks. They need to be set up. It has to be decided who runs and maintains them. And dynamic provisioning software needs to build identifiers by querying these registries. So the registries need an API for that.

The DANTE naming scheme needs a lot of coordination. E.g., a list of end site identifiers needs to be kept. Also, a list of project names needs to be kept. This will not work well in the GLIF community because of the many different sites and projects.

The Internet2 naming scheme needs no global registry. The sourcing organization picks its own DNS domain name and chooses a domain local part ..... (Tom?)

A problem with the Internet2 naming scheme is the overall length of the identifier because the DNS domain name can be quite long.

Using a hierarchical naming scheme that starts with a GOLE identifier needs a central repository with GOLE abbreviations. It makes more sense to use an existing registry like DNS, IATA or UNLOCODE.

Using a hierarchical naming scheme with the IATA airport of UN LOCODE as domain identifier solves the problem of long and variable domain names.

IATA codes are restricted to three characters, so they are short and have fixed length. However, they are only defined for airports and some railway stations and not quite applicable for the GLIF community.

UN LOCODEs consist of a two characters for the country and three characters for the city. So they have also a fixed short length. LOCODEs exist for far more places than IATA codes, but it is still difficult to find a proper LOCODE for some small villages.

When a flat naming scheme with guaranteed uniqueness is used, a global registry is needed. This is not the case for statistically unique identifiers. These can also have a fixed short length that can also be generated by dynamic provisioning software.

The DANTE and unique flat naming scheme need a global registry. Therefore we do not recommend these schemes.

The Internet2 DNS domain can be too large which causes trouble when storing it.

IATA codes do not exist for all places, so we do not recommend them either.

This leaves us with a hierarchical scheme starting with a UN LOCODE or a statistically unique flat namespace.

A hierarchical naming scheme has the advantage that each domain can decide for itself how it uses the domain local part. Although we recommend to restrict the length of it. A disadvantage is that all dynamic provisioning systems need to know how they must build the domain local part of global identifiers. Statistically unique flat naming schemes do not have this disadvantage. For dynamic provisioning systems they are the easiest to deal with.

## References

- [1] <http://en.wikipedia.org/wiki/un/locode>.