Global Lightpath Identifiers Proposal

Lars Fischer Tom Lehman Ronald van der Pol Thomas Tam

Draft 0.2
May, 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 Example 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

This scheme provides for a unique global identifier for each instance of a circuit reserved/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, ESNet, GEANT, and others.

2.3 Other 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 sourcing organization 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.

- Is a global registry needed?
- Is a central registry per domain needed?
- What is the maximum length of the identifier?
- Can the identifier be generated by provisioning software?

- Is the identifier unique or only statistically unique?
- What is the allowed character set?

An important question is where and how will global identifiers be used. Are they restricted purely as keys to other information or will they end up in interface descriptions, path or section traces, VLAN names, etc.?

In the latter case they need to be restricted in length. Otherwise they will be cut off and may not be unique anymore.

On the other hand, when global identifiers are only used as keys to other information about the lightpath, the total length should not be a problem.

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., lists of end site identifiers and project names need 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 uses its own DNS domain name and picks a number which is unique within the domain. The total length of the identifiers can be quite long when the DNS domain name is large.

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 (as in the Internet2 scheme), IATA or UN LOCODEs.

IATA codes are restricted to three characters, so they are short and have fixed length. However, they are only defined for aiports 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 also have 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 (which are typically the places where radio telescopes are built).

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

This leaves us with a three candidates:

- hierarchical scheme starting with a DNS domain name
- hierarchical scheme starting with a UN LOCODE
- statistically unique flat namespace

The hierarchical naming schemes have the advantage of containing some meaningful information about the lightpath. A DNS domain name is probably easier to pick than searching for a UN LOCODE. But it can make the total length of the identifier larger.

The scheme for the domain local part could be left up to each individual domain. But this makes the job a dynamic provisioning systems harder, because they need to handle each scheme. Therefore it seems wise to standardize on (statistically) unique numbers.

References

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