Semantic Models for Optical Hybrid Networks – Lightpaths Across Domain Boundaries

Paola GROSSO[1], Jeroen van der HAM[1], Freek DIJKSTRA[1], Cees de LAAT[1]
[1]*Universiteit van Amsterdam, Kruislaan 403, 1098SJ, Amsterdam, The Netherlands*
Tel: +31 20 525.7533, Fax: + 31 20 525.7460, Email: p.grosso@uva.nl

Abstract: Lightpaths are dedicated network connections in optical networks that provide guaranteed services to e-Science applications. When these paths extend beyond organizational boundaries the domains' provisioning systems need to inter-operate and exchange information on the underlying network topology and technology. We proposed the use of NDL - Network Description Language - as the vocabulary for this information exchange. NDL is based on Semantic Web and brings a new approach to exchange, distribution and storage of network topology information.

1. Introduction

Educational and research networks are now offering dedicated network connections to end-users. Such connections are called *lightpaths,* given they run as dedicated wavelengths (or *lambdas*) in the underlying optical networks. Lightpaths provide guaranteed network services to ever demanding e-Science applications that cannot use the traditional shared IP services. Networks architectures that support both lightpaths and routed IP paths are called hybrid networks [1].

Despite the fact that there are many examples of such setups around the world, we still do not see a common approach to the integration of lightpath services. Lightpaths do not cross seamlessly domain boundaries; most of the time they require manual configurations to go across organizations. One the biggest challenges to automated provisioning systems, and to the extension of lightpaths beyond organizational boundaries, is the lack of a common way to describe these connections and the optical networks that can deliver them. Lacking a common semantics network operators cannot easily integrate their tools with one another.

Now, Semantic Web is a way to organize data so that is contains 'meaning' that can be understood by machine and applications. It is also a way to correlate information that is distributed in different domains. Semantic Web data models (ontologies) describe relations among resources and their properties. The solutions offered by the Semantic Web map very well to our problem area. We believe that an ontology designed specifically for optical hybrid networks can provide the common vocabulary to allow lightpath-provisioning systems to exchange network topology information, and ultimately interoperate. These motivations are at the basis of the development of NDL – Network Description Language [2].

This paper outlines how ontologies can play an essential role in the inter-domain provisioning of lightpaths. It describes in details NDL and motivates the choices at the basis of the vocabulary definition. It shows some of the first results obtained using NDL for finding lightpaths across domains, and outlines the future research directions. We also indicate potential business and indutrial application of the language.

2. Interdomain lightpaths

All research networks in the world implementing a hybrid model use a provisioning system to provide lightpaths to their user base. These systems are either commercial products, such as the Nortel Networks DRAC in use in the SURFnet network in the Netherlands, or results of research efforts such as UCLP in the CAnet4 network in Canada, ARGUS in the Viola project in Germany, or DRAGON in the HOPI network in the USA. These systems differ in the way they setup the paths: this depends on the underlying network equipment they try to control. They also differ for the types of interfaces they provide to their users: in most cases they define Web Services APIs. In first instance they are not built with interoperability in mind.

Interoperability of provisioning systems is nonetheless essential for eScience applications, that more and more often involve international teams with data or computing resources distributed among domains. Similar to computing Grids where resources are geographically distributed and dynamically assigned to the applications depending on need and availability, inter-domain lightpaths form *lambda grids*. In a lambda grid lightpaths are assigned to applications for the duration of time needed to perform the computing tasks, are engineered according to the availability of the underlying network resources and span domains to reach the distributed resources.

The creation of inter-domain brokering systems that bridge the various local systems is the focus of several projects and international research initiatives. One among many others is the FP6-funded project called Phosphorus that addresses the technical challenges to enable on-demand e2e network services across multiple domains. Phosphorus will develop and disseminate procedures, toolkits and middleware to the EU NRENs and their users, such as Supercomputing centres and the wider European and worldwide scientific users.

Given this framework the best way to describe, categorize, store and distribute information between domain brokers and provisioning systems remains an open problem. Our answer is NDL – Network Description Language – based on Semantic Web and using RDF for the data description. Hybrid networks users and providers would rely on NDL to exchange information with a well-defined and unambiguous meaning.

3. Semantic approach to lightpath provisioning

While it would be possible to store the network information in databases in each domain or in a centralized location, we think that definition of ontologies and common RDF schemas brings additional functionalities and is ultimately more suited to solve our problem.

A Semantic Web approach [3] implies the definition of an ontology: every resource must be given a well-defined meaning. This is both an advantage to the schema author, who is forced to clearly define context and meaning for every single element, as well as for users, who may use this meaning to leverage other information on the Semantic Web.

Distributed information can be maintained in a structured way, allowing references between resources, both internally to a document and externally with the use of URIs. Network domains maintain full independence in publishing and updating their network information; while other network owners can still easily consume this information as long as it abides to the common ontology.

There is immediate guarantee of extensibility. A user can mix two ontologies in his application, even when neither ontology author was aware of the other schema. This means we can use lightpath ontologies in a concerted way with other separate ontologies to solve for example end-to-end data distribution. We will discuss this further in the 'Beyond lightpaths' section of this document. Finally, we can use the several existing tools that automatically 'make sense', process and consume semantically organized data.

3.1 RDF – Resource Description Framework

When we started to work on the semantic definition for hybrid networks we looked at how to describe our information. In the Semantic Web resources and properties can be described using RDF – Resource Description Framework. Instead of defining our semantic in flat XML files, we use the semantic richness of RDF to describe the relations between network resources. Furthermore using RDF as language for our vocabulary we can leverage all the existing Semantic Web tools that can parse, find and correlate information, speeding up the adoption rate of such a data model by the hybrid network research community.

RDF uses a triplet model, depicted in the following figure:

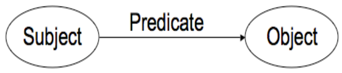

Figure 1- RDF triplet

In this figure the subject is a resource, the predicate is an associated property and the object is the property value. The object, or property value, can be also a resource itself, so that with RDF we can create chains of correlated information. Particularly useful for navigating between distributed RDF data is the property *isDefinedBy* that allows to reference information maintained elsewhere.

4. NDL schemas

NDL classes and properties are organized in five schemas:

- The *topology schema* that describes devices, interfaces and links and connections between them on a single layer;
- The *layer schema* that describes specific network technologies, and the relation between network layers;
- The *capability schema* that describes device capabilities;
- The *domain schema* that describes administrative domains, services within a domain, and how to give an abstracted view of the network in a domain;
- The *physical schema* that describes the physical aspects of network elements and the layout of devices with blades and chassis.

In the following two sections we give a more detailed overview of the topology schema and the layer schema. All other schemas are documented in the NDL web site [4].

4.1 NDL topology schema

The topology schema was the first schema we created. Its classes and properties describe the topology of a hybrid network, without detailed information on the technical aspects of the connections and their operating layer. The idea is that that through this lightweight schema we can provide an easy toolset for basic information exchange and path finding.

In Fig.2 we see the topology classes and properties:

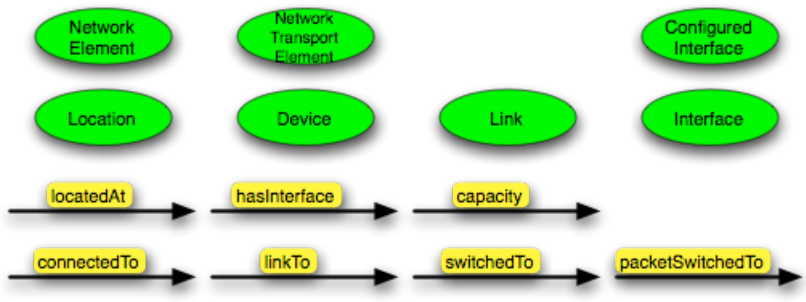

Figure 2 - NDL topology schema

Fig.3 shows the relation between the classes in this schema:

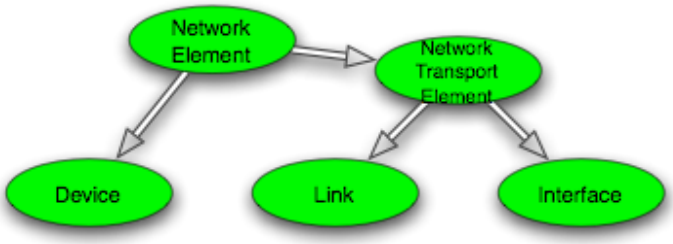

Figure 3 - NDL topology classes

A *Device* represents a physical or an abstract network element. In the topology schema we use the classes Interface and Link to create connections between devices. There are four properties to do this: *connectedTo* and *linkTo, switchedTo* and *packetSwitchedTo*.

The *connectedTo* property ties Interfaces together. It defines uni-directional connections between Interfaces, or from Interfaces to Links. All data sent out the subject Interface (the egress interface) is somehow received by to the object Interface (the ingress interface). To define a bi-directional connection with the connectedTo property both Interfaces should point to each other or both Interfaces pointing to the same Link. A connectedTo property always involves external connection, between devices, not connections within a device: for that we define the switchedTo property. The subject and object Interface must be on the same layer.

The *linkTo* is a subProperty of connectedTo. It defines that the subject and object Interface are directly connected to each other on their layer. There is no intermediate connection point forwarding the data.

The *switchedTo* property represents an internal uni-directional connection within a device. All data from the subject is forwarded to the object. A switchedTo property always involves internal connection, within devices, not connections between two devices. To define a bi-directional connection with the switchedTo property it should be defined in both directions.

The *packetSwitchedTo* property is a variant of the switchedTo property. When we define two switchedTo properties with the same Interface as object the result is that only one of the data streams is carried on the object Interface. The packetSwitchedTo properties on the other hand means that the data of the all subject Interfaces is somehow merged and forward on the object Interface.

The immediate applications of the topology schema are generation of network map and input to path finding systems. We will discuss this in more detail in the NDL applications section.

4.2 NDL layer schema

To actually provision a lightpath we need more information than the one contained in the topology schema. We need to define layers and relation between layers: this aids path finding in multilayer environments in which domains operate with different technologies. To facilitate this task we developed the NDL *layer schema.*

Fig. 4 shows the classes and properties in this schema.

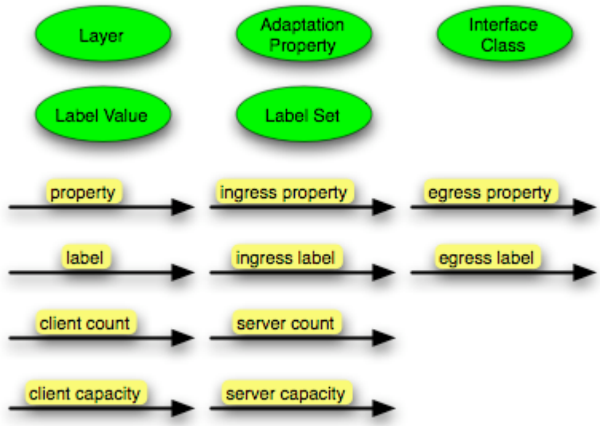

Figure 4 - NDL layer schema

A Layer is a specific encoding in network connection; most Layers have an associated Label Set that defines which channels are used to make switching decision in a device. Each Interface Class operates at a certain Layer. When data needs to move from one layer to another we use Adaptation. The concept of client and server refer to the Layers before and after the Adaptation. This way we can see if the requested path can be configured end-to-end. If the data needs to travel through devices that are on different layer, we need to check if the interfaces on these devices can adapt the data.

5. NDL applications

So far we have tested the NDL ontology in the Global Lambda Integrated Facility (GLIF) and in the SURFnet6 network in the Netherlands. GLIF is an international virtual organization of research networks, research consortia and institutions whose aim is to build a worldwide networking facility for scientific research. GLIF consists of a collection of optical exchange points, spread over the globe with numerous links across and between the continents. GLIF can provide scientists with international lightpaths, but these are created manually given the difficulty of integrating the various provisioning applications together.

Within GLIF we showed that NDL is suitable:

- to create network maps that give visual representations of connections between GLIF domains;
- to provide the input data to path finding tools.

5.1 Lightpath finding

During the international SuperComputing conference we showed for the first time the use of NDL as information base for path finding in the GLIF. This application was the result of cooperation with SARA, the Dutch Academic Super Computing Center, which contributed their path finding algorithms. All the GOLEs in the GLIF published an NDL description of their network. A crawler application gathered and correlated all this information navigating through the distributed data via the *isDefinedBy* property for interconnected interfaces. Path finding used this correlated information to determine which path were available between endpoints.

The path finding algorithms presented during SC|06 made use of the layer-agnostic topology schema. We are now looking at extending path finding algorithms to multi layer networks, where devices operate with different technologies. This use case is much more realistic in representing the actual implementations of lightpaths in intra and inter-domain environments. Multi-layer paths require the use of the layer schema and the capability schema.

5.2 Network abstractions

The GLIF demonstration is an example of centralized architecture: the information is distributed but the reasoning on the data is done in one place. A different approach is to de-centralize also the reasoning process and build a distributed architecture. Independent provisioning and path finding system exchange NDL information among them to calculate both the respective portions of the paths and the viability of path extension (and service extension) beyond the domain boundary. To facilitate this NDL needs to provide ontologies that do not describe the details of a network domain, but provide an abstracted view. The goal is to identify the minimal set of information necessary to extend lightpath at edges of domain. This is the focus of the under-development domain schema.

5.3 Policy information

In every hybrid network the domain's administrator and economic owners determine who is authorized to use their resources, which ones of the resources are available to this user and what kind of services this user can request for. Domain policies contain this information and AAA (Authentication Authorization and Accounting) servers take decision for access to the network based on such policies [5]. A provisioning system or network broker should not honor a request for a lightpath if this is not legitimate and allowed by the domain policies. A way to bridge the process of taking authorization decisions and provisioning the path is to include pointers to the policy repositories in the NDL files for a domain.

5.4 Topology requests

So far we have primarily targeted NDL primarily to hybrid network administrators and their provisioning system. We think that users and application requesting lightpaths can also benefit from such ontology. We imagine that a user looking for a path in a network would request the topology from the domain's broker, would reason on the possible paths and submit its request for a dedicated connection as a NDL file back to the broker. The ontology becomes not only the basis for full description of networks, but also for the interaction application-provisioning system. This is particularly useful when the desired topology is complex. We are using this approach in the StarPlane project where applications control a dedicated portion of the photonic network in SURFnet6.

6. NDL business potential

6.1 Standardization efforts

Within the research community there are several data models being developed to facilitate the management and operation of hybrid networks. NDL is one of them. All these efforts have the same goals, even if their approach slightly differ in the choice of the syntax language and the data categorization they choose. As all these data models are developed to provide a greater interoperability of provisioning tools, a standardized model that merges the best of these efforts is ideal.

The Open Grid Forum (OGF) [6] is a community of users, developers, and vendors leading the global standardization effort for grid computing. Within OGF the NDL developers are involved in the Network Markup Language Working Group (NML-WG) to cooperate with other researchers to the definition of a common vocabulary. The various provisioning tools currently deployed in the research networks would then use this standardized network description.

6.2 Business use cases

The provisioning of lightpaths is not simply a research endeavour. Network equipment vendors look at solutions to address this problem, so that they can deliver products that suit the needs of their customers. In this respect NDL, or similarly the standardized language that will result from the work in OGF, can be integrated in the commercial software used by research networks.

GMPLS [7] – Generalized Multi Protocol Label Switching – is a possible control plane for lightpath setup. It can automatically and in-band exchange network topology information between domains, and therefore between equipment of different vendors. The only problems to its widespread adoption are that commercial products need to implement it, and that administrators in each domain need to allow this mode of operation. Information about the capability and the configuration of the network devices that seamlessly travels between domains in an autonomous fashion does not always meet the criteria set by the usage policies for better control on the network assets.

We see in the research networks, and we would expect the same to happen in the commercial world, that an out-of-band approach to this exchange is often favoured. Web services in each domain answer to authorized requests with the information needed to setup the cross-domain lightpaths. NDL information is perfectly suited for such an exchange and we expect network vendors service the R&E community will consider exposing information using such model.

7. Beyond lightpaths

A use of inter-domain lightpaths emerging in the last period is the distribution of high-quality digital cinema and video content [8]. Initiatives such as CineGrid try ‘to build an interdisciplinary community that is focused on the research, development, and demonstration of networked collaborative tools to enable the production, use and exchange of very-high-quality digital media over photonic networks’. Lightpaths provide the appropriate level of service to efficiently deliver digital media, and as we illustrated in the preceding sections of this article NDL provides the basis for the provisioning of these connections. What we believe is the next step is the definition of ontologies that cover the whole end-to-end infrastructure: from the actual content being distributed to the Storage elements holding the data to the CPUs rendering the images and the display to visualize it. Figure 5 illustrates the concept of RDF representation of all this elements forming the overall infrastructure.

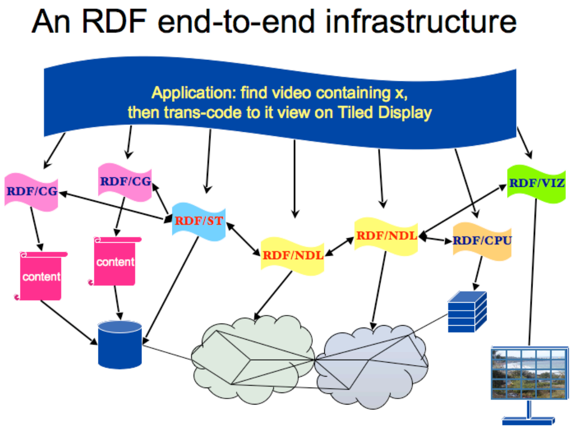

Figure 5 – 'End-to-end' ontologies

Our vision is that a media content locator will be able to consume the descriptions of all the architectural components that form the end-to-end infrastructure. From this information it wil build the optimal paths from the storage elements to the visualization displays, making the inter-domain lightpath provisioning a piece in the overall orchestrated effort.

8. Conclusions

NDL – Network Description Language – is a Semantic Web data model that describes optical networks. It is specifically designed to aid provisioning applications to find and construct lightpaths between domains. Such ontology provides the common and shared vocabulary that has so far been missing. As more domains adopt this vocabulary NDL will become one of the vehicle for topology information transfers between domains. This has already been proved in the GLIF and SURFnet6 networks. NDL extensions will furthermore allows references to authorization policies and integration with the ontologies developed to cover the end-to-end application's operational space.

9. Acknowledgements

We thank our colleagues Andree Toonk and Ronald van der Pol from SARA and Franco Travostino from eBay (formerly NORTEL) for their collaboration on this work. Parts of this work are supported by TNO, the EU project Phosphorus, the GigaPort project and by SURFnet.

References

[1] Cees de Laat, Erik Radius, Steven Wallace In: "The rationale of the current optical networking initiatives", Future Generation Computer Systems, Vol.22, Issue 6, pp. 999-1008, 2003
[2] Jeroen van der Ham, Freek Dijkstra, Franco Travostino, Hubertus Andree and Cees de Laat In: "Using RDF to Describe Networks", Future Generation Computer Systems, Vol. 22, Issue 8, pp. 908-914, 2006
[3] Nigel Shadbolt, Tim Berners-Lee, and Wendy Hall In: "The Semantic Web Revisited", IEEE Intelligent Systems, Vol. 21, Issue 3, pp. 96 -101, 2006
[4] NDL – Network Description Language, URL: http://www.science.uva.nl/research/sne/ndl/
[5] Leon Gommans, Freek Dijkstra, Cees de Laat, Arie Taal, Alfred Wan, Bas van Oudenaarde, Tal Lavian, Inder Monga, Franco Travostino, In: "Applications Drive Secure Lightpath Creation across Heterogeneous Domains", IEEE Communications Magazine, vol. 44, no. 3, March 2006
[6] Open Grid Forum:, URL: http://www.ogf.org
[7] Adrian Farrel and Igor Bryskin, In: "GMPLS Architecture and Applications", Publ. Morgan Kaufman, 2006
[8] Cees de Laat, Paola Grosso , In "Lambda Grid developments, History - Present – Future" Lighting the Blue Touchpaper for UK e-Science - Closing Conference of ESLEA Project, Proceedings of Science, June 2007